



ExtremeXOS Command Reference Guide

Software Version 12.0

Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, California 95051
(888) 257-3000
(408) 579-2800
<http://www.extremenetworks.com>

Published: 2007
Part number: 100261-00 Rev 01



AccessAdapt, Alpine, BlackDiamond, ESRP, Ethernet Everywhere, Extreme Enabled, Extreme Ethernet Everywhere, Extreme Networks, Extreme Standby Router Protocol, Extreme Turbodriven, Extreme Velocity, ExtremeWare, ExtremeWorks, ExtremeXOS, the Go Purple Extreme Solution, Sentrant, ServiceWatch, ScreenPlay, Summit, SummitStack, Unified Access Architecture, Unified Access RF Manager, UniStack, Universal Port, the Extreme Networks logo, the Alpine logo, the BlackDiamond logo, the Extreme Turbodriven logo, the Summit logos, the Powered by ExtremeXOS logo, and the Color Purple, among others, are trademarks or registered trademarks of Extreme Networks, Inc. or its subsidiaries in the United States and/or other countries.

Specifications are subject to change without notice.

Adobe, Flash, and Macromedia are registered trademarks of Adobe Systems Incorporated in the U.S. and/or other countries. AutoCell is a trademark of AutoCell. Avaya is a trademark of Avaya, Inc. Internet Explorer is a registered trademark of Microsoft Corporation. Merit is a registered trademark of Merit Network, Inc. Mozilla Firefox is a registered trademark of the Mozilla Foundation. sFlow is a registered trademark of sFlow.org. Solaris and Java are trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All other trademarks, registered trademarks, and service marks are property of their respective owners.

© 2007 Extreme Networks, Inc. All Rights Reserved.

Contents

Preface.....	39
Chapter 1: Command Reference Overview.....	43
Chapter 2: Commands for Accessing the Switch.....	55
clear account logout	56
clear license-info	57
clear session.....	58
configure account	59
configure account encrypted	61
configure account password-policy char-validation.....	63
configure account password-policy history	65
configure account password-policy logout-on-login-failures	67
configure account password-policy max-age	69
configure account password-policy min-length	71
configure banner.....	73
configure cli max-sessions	74
configure cli max-failed-logins.....	75
configure dns-client add	76
configure dns-client default-domain.....	78
configure dns-client delete.....	79
configure failsafe-account.....	81
configure idletimeout	83
configure safe-default-script.....	84
configure time	86
configure timezone.....	87
create account	91
delete account	93
disable cli space-completion	95
disable clipaging.....	96
disable idletimeout.....	97
enable cli space-completion	98
enable clipaging.....	99
enable idletimeout	100
enable license software.....	101
enable license file	104
history.....	106
ping.....	107
reboot	109
show accounts	111
show accounts password-policy.....	113

show banner	115
show dns-client.....	116
show licenses	117
show switch.....	119
traceroute.....	123

Chapter 3: Commands for Managing the Switch 125

configure node priority.....	128
configure power supply	130
configure snmp access-profile	132
configure snmp add community.....	134
configure snmp add trapreceiver.....	136
configure snmp delete community	138
configure snmp delete trapreceiver	140
configure snmp sysContact.....	141
configure snmp sysLocation	142
configure snmp sysName	143
configure snmpv3 add access.....	144
configure snmpv3 add community	146
configure snmpv3 add filter	148
configure snmpv3 add filter-profile	150
configure snmpv3 add group user	151
configure snmpv3 add mib-view	153
configure snmpv3 add notify	155
configure snmpv3 add target-addr	156
configure snmpv3 add target-params	158
configure snmpv3 add user	160
configure snmpv3 add user clone-from.....	162
configure snmpv3 delete access	163
configure snmpv3 delete community.....	165
configure snmpv3 delete filter	166
configure snmpv3 delete filter-profile.....	167
configure snmpv3 delete group user.....	168
configure snmpv3 delete mib-view	170
configure snmpv3 delete notify.....	172
configure snmpv3 delete target-addr	173
configure snmpv3 delete target-params.....	174
configure snmpv3 delete user.....	175
configure snmpv3 engine-boots	177
configure snmpv3 engine-id	178
configure sntp-client	179
configure sntp-client update-interval.....	181
configure telnet access-profile.....	182
configure telnet port.....	184
configure telnet vr	186
disable dhcp vlan	188

disable snmp access	189
disable snmp traps	190
disable snmp-client	191
disable telnet.....	192
disable watchdog	193
enable dhcp vlan.....	194
enable snmp access	195
enable snmp traps.....	197
enable snmp-client	198
enable telnet	199
enable watchdog	201
exit	202
logout	203
quit	204
show dhcp-client state.....	205
show checkpoint-data	206
show management	208
show node	211
show odometers	213
show power	216
show power budget.....	220
show power controller	223
show session	225
show snmpv3 access	227
show snmpv3 context	230
show snmpv3 counters	231
show snmpv3 engine-info.....	233
show snmpv3 filter	234
show snmpv3 filter-profile.....	236
show snmpv3 group.....	238
show snmpv3 mib-view	241
show snmpv3 notify.....	243
show snmpv3 target-addr	244
show snmpv3 extreme-target-addr-ext	246
show snmpv3 target-params.....	247
show snmpv3 user.....	249
show snmp-client.....	251
telnet.....	253
telnet msm	255
telnet slot.....	257
tftp	259
tftp get.....	263
tftp put	266
Chapter 4: Commands for Managing the ExtremeXOS Software.....	269
clear cpu-monitoring	270

cp.....	272
disable cpu-monitoring	275
disable xml-mode.....	276
enable cpu-monitoring.....	277
enable xml-mode	278
ls.....	279
mv.....	282
restart process	285
rm	288
show cpu-monitoring	291
show heartbeat process.....	295
show memory.....	298
show memory process.....	303
show process	306
start process.....	313
terminate process.....	315
Chapter 5: Commands for Configuring Slots and Ports on a Switch.....	319
clear counters ports.....	321
clear lacp counters.....	322
clear counters edp.....	323
clear slot	324
configure edp advertisement-interval.....	325
configure ip-mtu vlan	326
configure jumbo-frame-size	328
configure lacp member-port priority	330
configure mirror add anomaly.....	331
configure mirroring add	332
configure mirroring delete	337
configure mirroring mode.....	338
configure ports auto off.....	339
configure ports auto on	341
configure ports auto-polarity.....	342
configure ports display-string	344
configure ports preferred-medium.....	345
configure ports rate-limit flood	347
configure ports redundant	348
configure ports wan-phy clocking.....	350
configure ports wan-phy framing.....	351
configure ports wan-phy loopback.....	352
configure ports wan-phy trace-path	353
configure ports wan-phy trace-section	354
configure sharing add ports.....	355
configure sharing address-based.....	358
configure sharing delete ports	360
configure sharing lacp activity-mode	361

configure sharing lacp defaulted-state-action	363
configure sharing lacp system-priority	365
configure sharing lacp timeout	367
configure slot module	368
configure switch ports initial-mode disabled	370
configure switch ports initial-mode enabled	371
disable edp ports	372
disable jumbo-frame ports	373
disable learning port.....	375
disable mirroring.....	376
disable port	377
disable sharing	378
disable slot.....	379
disable smartredundancy	381
disable snmp traps port-up-down ports.....	382
enable edp ports	383
enable jumbo-frame ports	385
enable learning port	387
enable mirroring to port	388
enable port.....	392
enable sharing grouping.....	393
enable slot	398
enable smartredundancy	400
enable snmp traps port-up-down ports.....	401
restart ports.....	402
run failover	403
run msm-failover	405
show edp.....	406
show lacp.....	408
show lacp counters.....	410
show lacp lag.....	412
show lacp member-port	417
show mirroring.....	420
show ports anomaly.....	422
show ports collisions	424
show ports configuration	426
show ports information	429
show ports packet	435
show ports redundant	437
show ports sharing	438
show ports utilization.....	440
show ports wan-phy configuration	445
show ports wan-phy errors.....	447
show ports wan-phy events.....	449
show ports wan-phy overhead	451
show slot.....	453

unconfigure ports display string	462
unconfigure ports redundant	463
unconfigure ports wan-phy	464
unconfigure slot	465
Chapter 6: Universal Port Management Commands.....	467
configure upm event.....	468
configure upm timer after	469
configure upm timer at	470
configure upm timer profile.....	472
create upm profile.....	473
create upm timer	474
delete upm profile.....	475
delete upm timer	476
disable upm profile	477
edit upm profile	478
enable upm profile	479
show upm events	480
show upm history	482
show upm profile	484
show upm timers.....	486
unconfigure upm event	488
unconfigure upm timer	489
Chapter 7: CLI Scripting Commands.....	491
configure cli mode	492
configure cli mode scripting.....	493
delete var key	494
disable cli scripting.....	495
ELSE	496
enable cli scripting.....	497
ENDIF.....	498
ENDWHILE	499
IF ... THEN.....	500
load var key	502
save var key	503
set var.....	504
show var.....	505
WHILE ... DO.....	507
Chapter 8: Commands for Connectivity Fault Management	509
configure cfm domain add association integer.....	510
configure cfm domain add association string.....	511
configure cfm domain add association vpn-id oui index.....	512
configure cfm domain association add end-point.....	513
configure cfm domain association add filter-function	515
configure cfm domain association add intermediate-point	517

configure cfm domain association delete end-point	519
configure cfm domain association delete filter-function	520
configure cfm domain association delete intermediate-point.....	521
configure cfm domain association end-point mepid	522
configure cfm domain association end-point transmit-interval.....	523
configure cfm domain delete association	525
configure cfm domain ma-level.....	526
create cfm domain dns ma-level	527
create cfm domain iso3166 ma-level	528
create cfm domain m1400 ma-level	529
create cfm domain mac ma-level	530
create cfm domain vpn-id oui index ma-level	531
delete cfm domain	532
ping mac port	533
show cfm	535
show cfm detail	537
show cfm statistics.....	539
traceroute mac port.....	541
Chapter 9: Commands for Configuring LLDP	543
configure lldp med fast-start repeat-count	544
configure lldp ports management-address.....	545
configure lldp ports port-description.....	546
configure lldp ports system-capabilities.....	547
configure lldp ports system-description	548
configure lldp ports system-name	549
configure lldp ports vendor-specific avaya-extreme call-server.....	550
configure lldp ports vendor-specific avaya-extreme dot1q-framing.....	551
configure lldp ports vendor-specific avaya-extreme file-server	553
configure lldp ports vendor-specific avaya-extreme poe-conservation-request.....	554
configure lldp ports vendor-specific dot1 port-vlan-ID	556
configure lldp ports vendor-specific dot1 port-protocol-vlan-ID	557
configure lldp ports vendor-specific dot1 vlan-name	559
configure lldp ports vendor-specific dot3 link-aggregation	561
configure lldp ports vendor-specific dot3 mac-phy.....	562
configure lldp ports vendor-specific dot3 max-frame-size	563
configure lldp ports vendor-specific dot3 power-via-mdi	564
configure lldp ports vendor-specific med capabilities.....	566
configure lldp ports vendor-specific med location-identification	568
configure lldp ports vendor-specific med policy application	570
configure lldp ports vendor-specific med power-via-mdi	573
configure lldp reinitialize-delay.....	575
configure lldp snmp-notification-interval	576
configure lldp transmit-delay.....	577
configure lldp transmit-hold	578
configure lldp transmit-interval.....	579

disable lldp ports	580
disable snmp traps lldp	581
disable snmp traps lldp-med	582
enable lldp ports	583
enable snmp traps lldp	584
enable snmp traps lldp-med	585
show lldp	586
show lldp neighbors.....	589
show lldp statistics.....	592
unconfigure lldp.....	594
Chapter 10: Power Over Ethernet Commands.....	595
clear inline-power stats ports.....	596
configure inline-power budget	598
configure inline-power disconnect-precedence	600
configure inline-power label ports	602
configure inline-power operator-limit ports.....	604
configure inline-power priority ports	606
configure inline-power usage-threshold.....	607
disable inline-power	608
disable inline-power legacy	610
disable inline-power legacy slot	611
disable inline-power ports	612
disable inline-power slot	614
enable inline-power	615
enable inline-power legacy	617
enable inline-power legacy slot	618
enable inline-power ports.....	620
enable inline-power slot.....	622
reset inline-power ports	624
show inline-power	625
show inline-power configuration ports.....	627
show inline-power info ports.....	629
show inline-power slot	632
show inline-power stats.....	634
show inline-power stats ports.....	635
show inline-power stats slot.....	637
unconfigure inline-power budget slot.....	639
unconfigure inline-power disconnect-precedence.....	640
unconfigure inline-power operator-limit ports	641
unconfigure inline-power priority ports	642
unconfigure inline-power usage-threshold	643
Chapter 11: Commands for Status Monitoring and Statistics	645
clear counters.....	649
clear elsm ports auto-restart.....	650

clear elsm ports counters	651
clear log	652
clear log counters	653
clear sys-recovery-level	655
configure elsm ports hellotime	657
configure elsm ports hold-threshold	659
configure elsm ports uptime-threshold	661
configure log display	662
configure log filter events	664
configure log filter events match	667
configure log target filter	671
configure log target format	674
configure log target match	678
configure log target severity	680
configure log target syslog	682
configure ports monitor vlan	684
configure sflow agent ipaddress	686
configure sflow collector ipaddress	687
configure sflow max-cpu-sample-limit	689
configure sflow poll-interval	690
configure sflow ports sample-rate	691
configure sflow sample-rate	693
configure sys-health-check all level	694
configure sys-health-check interval	697
configure sys-recovery-level	700
configure sys-recovery-level slot	702
configure sys-recovery-level switch	709
configure syslog add	711
configure syslog delete	713
create log filter	715
delete log filter	717
disable cli-config-logging	718
disable elsm ports	719
disable elsm ports auto-restart	720
disable log debug-mode	722
disable log display	723
disable log target	724
disable rmon	726
disable sflow	727
disable sflow ports	728
disable sys-health-check	729
disable syslog	731
enable cli-config-logging	732
enable elsm ports	733
enable elsm ports auto-restart	735
enable log debug-mode	737

enable log display	738
enable log target	740
enable rmon	742
enable sflow	744
enable sflow ports	746
enable sys-health-check.....	748
enable syslog	750
show elsm	751
show elsm ports	754
show fans.....	759
show log.....	763
show log components	767
show log configuration	771
show log configuration filter	773
show log configuration target.....	775
show log counters	778
show log events.....	780
show ports rxerrors	782
show ports statistics	785
show ports txerrors	788
show ports vlan statistics	791
show rmon memory	793
show sflow configuration	796
show sflow statistics.....	798
show temperature.....	800
show version	803
show vlan statistics	808
unconfigure log filter	809
unconfigure log target format	810
unconfigure ports monitor vlan	812
unconfigure sflow.....	813
unconfigure sflow agent.....	814
unconfigure sflow collector.....	815
unconfigure sflow ports.....	816
upload log	817
Chapter 12: VLAN Commands.....	819
clear mac-binding	820
configure bvlan add ports tagged	821
configure bvlan add svlan	822
configure bvlan delete ports	823
configure bvlan delete svlan	824
configure bvlan ethertype.....	825
configure bvlan tag.....	826
configure protocol add	827
configure protocol delete	829

configure svlan add ports	830
configure svlan delete ports.....	831
configure svlan isid	832
configure svlan tag	833
configure vlan add ports	834
configure vlan delete ports	836
configure vlan ipaddress	837
configure vlan name	840
configure vlan protocol	841
configure vlan tag	843
create bvlan.....	844
create protocol.....	845
create svlan	846
create vlan	847
delete bvlan.....	849
delete protocol.....	850
delete svlan	851
delete vlan	852
disable dot1p examination inner-tag ports	853
disable loopback-mode vlan	854
disable vlan	855
enable dot1p examination inner-tag port	856
enable loopback-mode vlan	857
enable vlan.....	858
show bvlan	859
show mac-binding.....	861
show protocol	863
show svlan.....	864
show vlan	866
unconfigure vlan ipaddress.....	870
Chapter 13: VMAN Commands.....	873
configure flood-group extended add	874
configure flood-group extended delete.....	875
configure flood-group virtual add	876
configure flood-group virtual delete.....	878
configure port ethertype.....	880
create flood-group extended	881
create flood-group virtual	882
configure vman add ports.....	883
configure vman delete ports	886
configure vman ethertype.....	887
configure vman tag.....	889
create vman.....	890
delete vman.....	892
show vman	893

unconfigure vman ethertype	896
Chapter 14: FDB Commands	897
clear fdb	898
configure fdb agingtime	899
create fdbentry vlan ports	900
delete fdbentry	902
disable flooding port	903
enable flooding port	906
show fdb	908
Chapter 15: Commands for Virtual Routers	911
create virtual-router	912
configure vr add protocol	913
configure vr add ports	914
configure vr delete protocol	915
configure vr delete ports	916
delete virtual-router	917
show virtual-router	918
virtual-router	920
Chapter 16: Policy Manager Commands	923
check policy	924
check policy attribute	925
edit policy	927
refresh policy	929
show policy	931
Chapter 17: ACL Commands	933
clear access-list counter	934
clear access-list meter	936
configure access-list	937
configure access-list add	939
configure access-list delete	941
configure access-list zone	942
create access-list	944
create access-list zone	946
delete access-list	947
delete access-list zone	948
disable access-list permit to-cpu	949
disable access-list refresh blackhole	950
enable access-list permit to-cpu	951
enable access-list refresh blackhole	952
show access-list	953
show access-list configuration	955
show access-list counter	956
show access-list dynamic	958

show access-list dynamic counter	959
show access-list dynamic rule	960
show access-list interface	962
show access-list meter.....	965
show access-list usage acl-mask	966
show access-list usage acl-range.....	967
show access-list usage acl-rule	969
show access-list usage acl-slice	971
unconfigure access-list	973
Chapter 18: QoS Commands	975
clear counters traffic queue.....	977
configure diffserv examination code-point.....	978
configure diffserv replacement	980
configure dot1p type	982
configure meter	984
configure meter committed-rate peak-rate	986
configure meter peak-rate	988
configure ports qosprofile.....	989
configure ports rate-limit egress.....	991
configure ports rate-limit egress aggregate.....	993
configure ports rate-limit map cos high	995
configure ports rate-limit packet byte-adjustment.....	998
configure qosprofile ingress ports.....	1000
configure qosprofile ports.....	1002
configure qosprofile weight	1005
configure qosscheduler	1007
configure traffic ingress queue add egress queue	1008
configure traffic ingress queue delete egress queue	1010
configure traffic mode	1011
configure traffic queue aggregate-meter	1012
configure traffic queue cos.....	1013
configure vlan qosprofile.....	1015
configure vman rate-limit flood-traffic designated-port	1016
create meter	1018
create qosprofile	1019
create traffic queue allow-egress-shaping	1021
create traffic queue egress.....	1023
create traffic queue ingress-only	1025
delete meter	1026
delete qosprofile	1027
delete traffic queue	1028
disable diffserv examination port	1029
disable diffserv replacement port.....	1030
disable dot1p examination ports.....	1031
disable dot1p replacement ports.....	1032

enable diffserv examination port	1033
enable diffserv replacement ports	1034
enable dot1p examination ports	1036
enable dot1p replacement ports	1037
show access-list meter	1039
show diffserv	1040
show dot1p	1042
show meter	1044
show ports qosmonitor	1047
show qosprofile	1049
show traffic mode	1052
show traffic queue	1053
show traffic queue statistics	1056
show traffic queue utilization	1059
unconfigure diffserv	1061
unconfigure ports aggregate-bandwidth	1062
unconfigure qosprofile	1063
unconfigure vman rate-limit flood-traffic designated-port	1066
Chapter 19: Security Commands	1069
clear ip-security anomaly-protection notify cache	1070
clear ip-security dhcp-snooping entries	1071
clear ip-security source-ip-lockdown entries ports	1072
clear vlan dhcp-address-allocation	1073
configure dos-protect acl-expire	1074
configure dos-protect interval	1075
configure dos-protect trusted ports	1076
configure dos-protect type I3-protect alert-threshold	1077
configure dos-protect type I3-protect notify-threshold	1078
configure ip-security anomaly-protection icmp ipv4-max-size	1079
configure ip-security anomaly-protection icmp ipv6-max-size	1080
configure ip-security anomaly-protection notify cache	1081
configure ip-security anomaly-protection notify rate limit	1082
configure ip-security anomaly-protection notify rate window	1083
configure ip-security anomaly-protection notify trigger off	1084
configure ip-security anomaly-protection notify trigger on	1085
configure ip-security anomaly-protection tcp	1086
configure mac-lockdown-timeout ports aging-time	1087
configure ports limit-learning lock-learning	1088
configure radius server	1091
configure radius shared-secret	1093

configure radius timeout	1095
configure radius-accounting server	1096
configure radius-accounting shared-secret	1098
configure radius-accounting timeout	1100
configure ssh2 key	1101
configure sshd2 user-key add user	1103
configure sshd2 user-key delete user	1104
configure ssl certificate pregenerated	1105
configure ssl certificate privkeylen	1107
configure ssl privkey pregenerated	1109
configure tacacs server	1111
configure tacacs shared-secret	1113
configure tacacs timeout	1114
configure tacacs-accounting server	1115
configure tacacs-accounting shared-secret	1117
configure tacacs-accounting timeout	1118
configure trusted-ports trust-for dhcp-server	1119
configure trusted-servers vlan add server trust-for dhcp-server	1121
configure trusted-servers vlan delete server trust-for dhcp-server	1123
configure vlan dhcp-address-range	1124
configure vlan dhcp-lease-timer	1125
configure vlan dhcp-options	1126
create sshd2 key-file	1127
create sshd2 user-key	1128
delete sshd2 user-key	1129
disable dhcp ports vlan	1130
disable dos-protect	1131
disable iparp gratuitous protect	1132
disable ip-security anomaly-protection	1133
disable ip-security anomaly-protection ip	1134
disable ip-security anomaly-protection l4port	1135
disable ip-security anomaly-protection tcp flags	1136
disable ip-security anomaly-protection tcp fragment	1137
disable ip-security anomaly-protection icmp	1138
disable ip-security anomaly-protection notify	1139
disable ip-security arp gratuitous-protection vlan	1140
disable ip-security arp learning learn-from-arp vlan ports	1141
disable ip-security arp learning learn-from-dhcp vlan ports	1143
disable ip-security arp validation vlan	1144
disable ip-security dhcp-snooping	1145
disable ip-security source-ip-lockdown ports	1146
disable mac-lockdown-timeout ports	1147
disable radius	1148
disable radius-accounting	1149
disable ssh2	1150
disable tacacs	1151

disable tacacs-accounting	1152
disable tacacs-authorization	1153
disable web http	1154
disable web https	1155
download ssl certificate	1156
download ssl privkey	1158
enable dhcp ports vlan	1160
enable dos-protect	1161
enable dos-protect simulated	1162
enable iparp gratuitous protect	1163
enable ip-security anomaly-protection	1164
enable ip-security anomaly-protection icmp	1165
enable ip-security anomaly-protection ip	1166
enable ip-security anomaly-protection l4port	1167
enable ip-security anomaly-protection notify	1168
enable ip-security anomaly-protection tcp flags	1169
enable ip-security anomaly-protection tcp fragment	1170
enable ip-security arp gratuitous-protection vlan	1171
enable ip-security arp learning learn-from-arp vlan ports	1173
enable ip-security arp learning learn-from-dhcp vlan ports	1175
enable ip-security arp validation vlan violation-action	1177
enable ip-security dhcp-snooping vlan ports violation-action	1179
enable ip-security source-ip-lockdown ports	1181
enable mac-lockdown-timeout ports	1183
enable radius	1184
enable radius-accounting	1186
enable ssh2	1188
enable tacacs	1190
enable tacacs-accounting	1191
enable tacacs-authorization	1192
enable web http	1193
enable web https	1194
scp2	1196
show dhcp-server	1199
show dos-protect	1200
show ip-security anomaly-protection notify cache ports	1202
show ip-security arp gratuitous-protection	1203
show ip-security arp learning vlan	1204
show ip-security arp validation vlan	1206
show ip-security dhcp-snooping entries vlan	1207
show ip-security dhcp-snooping violations vlan	1208
show ip-security dhcp-snooping vlan	1209
show ip-security source-ip-lockdown	1211
show mac-lockdown-timeout fdb ports	1212
show mac-lockdown-timeout ports	1214
show radius	1216

show radius-accounting	1219
show sshd2 user-key	1221
show ssl	1222
show tacacs	1224
show tacacs-accounting	1226
show vlan dhcp-address-allocation	1228
show vlan dhcp-config	1229
show vlan security	1230
ssh2	1231
unconfigure radius	1234
unconfigure radius-accounting	1236
unconfigure tacacs	1238
unconfigure tacacs-accounting	1239
unconfigure trusted-ports trust-for dhcp-server	1240
unconfigure vlan dhcp	1241
unconfigure vlan dhcp-address-range	1242
unconfigure vlan dhcp-options	1243
Chapter 20: Network Login Commands	1245
clear netlogin state	1246
clear netlogin state mac-address	1247
configure netlogin add mac-list	1248
configure netlogin agingtime	1250
configure netlogin authentication database-order	1251
configure netlogin banner	1252
configure netlogin base-url	1253
configure netlogin delete mac-list	1254
configure netlogin dot1x eapol-transmit-version	1255
configure netlogin dot1x guest-vlan	1256
configure netlogin dot1x timers	1258
configure netlogin dynamic-vlan	1260
configure netlogin dynamic-vlan uplink-ports	1262
configure netlogin local-user	1264
configure netlogin move-fail-action	1266
configure netlogin ports mode	1267
configure netlogin ports no-restart	1270
configure netlogin ports restart	1271
configure netlogin redirect-page	1272
configure netlogin vlan	1273
configure vlan netlogin-lease-timer	1274
create netlogin local-user	1275
delete netlogin local-user	1278
disable netlogin	1279
disable netlogin dot1x guest-vlan ports	1280
disable netlogin logout-privilege	1281
disable netlogin ports	1282

disable netlogin session-refresh	1283
enable netlogin	1284
enable netlogin dot1x guest-vlan ports	1285
enable netlogin logout-privilege	1287
enable netlogin ports	1288
enable netlogin session-refresh	1289
show banner netlogin	1290
show netlogin	1291
show netlogin guest-vlan	1297
show netlogin local-users	1298
show netlogin mac-list	1300
unconfigure netlogin authentication database-order	1301
unconfigure netlogin dot1x guest-vlan	1302
unconfigure netlogin vlan	1303
Chapter 21: CLEAR-Flow Commands	1305
disable clear-flow	1306
enable clear-flow	1307
show clear-flow	1308
show clear-flow acl-modified	1309
show clear-flow rule	1310
show clear-flow rule-all	1312
show clear-flow rule-triggered	1314
Chapter 22: SummitStack Commands	1315
configure stacking alternate-ip-address	1316
configure stacking easy-setup	1318
configure stacking license-level	1321
configure stacking mac-address	1323
configure stacking master-capability	1325
configure stacking priority	1327
configure stacking redundancy	1329
configure stacking slot-number	1331
disable stacking	1333
enable stacking	1335
show power (SummitStack Only)	1337
show stacking	1339
show stacking configuration	1342
show stacking detail	1344
show stacking stack-ports	1346
synchronize stacking	1348
unconfigure stacking	1350
unconfigure stacking alternate-ip-address	1352
unconfigure stacking license-level	1353
Chapter 23: EAPS Commands	1355
clear eaps counters	1357

configure eaps add control vlan	1358
configure eaps add protected vlan.....	1360
configure eaps config-warnings off	1362
configure eaps config-warnings on	1364
configure eaps delete control vlan.....	1365
configure eaps delete protected vlan	1366
configure eaps failtime	1368
configure eaps failtime expiry-action	1370
configure eaps fast-convergence	1371
configure eaps hellotime.....	1373
configure eaps mode	1374
configure eaps name	1376
configure eaps port	1377
configure eaps shared-port link-id.....	1379
configure eaps shared-port mode	1380
configure eaps shared-port segment-timeout expiry-action.....	1381
create eaps.....	1383
create eaps shared-port	1384
delete eaps.....	1385
delete eaps shared-port	1386
disable eaps	1387
enable eaps	1389
show eaps	1391
show eaps counters	1395
show eaps counters shared-port	1400
show eaps shared-port	1405
show vlan eaps	1410
unconfigure eaps shared-port link-id	1412
unconfigure eaps shared-port mode	1413
unconfigure eaps port.....	1414
Chapter 24: STP Commands	1417
clear counters stp	1422
configure mstp format	1423
configure mstp region	1424
configure mstp revision.....	1426
configure stpd add vlan	1427
configure stpd default-encapsulation.....	1430
configure stpd delete vlan	1432
configure stpd forwarddelay	1433
configure stpd hellotime	1434
configure stpd maxage.....	1435
configure stpd max-hop-count	1436
configure stpd mode.....	1438
configure stpd ports cost.....	1440
configure stpd ports edge-safeguard disable.....	1442

configure stpd ports edge-safeguard enable	1444
configure stpd ports link-type	1446
configure stpd ports mode	1448
configure stpd ports port-priority	1450
configure stpd ports priority	1452
configure stpd priority	1454
configure stpd tag	1456
configure vlan add ports stpd	1458
create stpd	1461
delete stpd	1463
disable stpd	1464
disable stpd auto-bind	1465
disable stpd ports	1466
disable stpd rapid-root-failover	1467
enable stpd	1468
enable stpd auto-bind	1469
enable stpd ports	1472
enable stpd rapid-root-failover	1473
show stpd	1474
show stpd ports	1478
show vlan stpd	1481
unconfigure mstp region	1483
unconfigure stpd	1484
unconfigure stpd ports link-type	1485
Chapter 25: ESRP Commands	1487
clear esrp counters	1490
clear esrp neighbor	1492
clear esrp sticky	1493
configure esrp add elrp-poll ports	1494
configure esrp add master	1495
configure esrp add member	1496
configure esrp add track-environment	1497
configure esrp add track-iproute	1498
configure esrp add track-ping	1499
configure esrp add track-vlan	1500
configure esrp aware add selective-forward-ports	1501
configure esrp aware delete selective-forward-ports	1502
configure esrp delete elrp-poll ports	1503
configure esrp delete master	1504
configure esrp delete member	1505
configure esrp delete track-environment	1506
configure esrp delete track-iproute	1507
configure esrp delete track-ping	1508
configure esrp delete track-vlan	1509
configure esrp domain-id	1510

configure esrp election-policy	1511
configure esrp elrp-master-poll disable	1515
configure esrp elrp-master-poll enable	1516
configure esrp elrp-premaster-poll disable	1518
configure esrp elrp-premaster-poll enable	1519
configure esrp group	1521
configure esrp mode	1523
configure esrp name	1525
configure esrp ports mode	1526
configure esrp ports no-restart	1527
configure esrp ports restart	1528
configure esrp ports weight	1529
configure esrp priority	1530
configure esrp timer hello	1531
configure esrp timer neighbor	1532
configure esrp timer neutral	1534
configure esrp timer premaster	1535
configure esrp timer restart	1536
create esrp	1537
delete esrp	1539
disable esrp	1540
enable esrp	1541
show esrp	1542
show esrp aware	1545
show esrp counters	1547
Chapter 26: VRRP Commands	1549
clear counters vrrp	1551
configure vrrp vlan vrid add ipaddress	1552
configure vrrp vlan vrid add track-iproute	1553
configure vrrp vlan vrid add track-ping	1554
configure vrrp vlan vrid add track-vlan	1555
configure vrrp vlan vrid advertisement-interval	1556
configure vrrp vlan vrid authentication	1558
configure vrrp vlan vrid delete ipaddress	1559
configure vrrp vlan vrid delete track-iproute	1560
configure vrrp vlan vrid delete track-ping	1561
configure vrrp vlan vrid delete track-vlan	1562
configure vrrp vlan vrid dont-preempt	1563
configure vrrp vlan vrid preempt	1564
configure vrrp vlan vrid priority	1565
create vrrp vlan vrid	1566
delete vrrp vlan vrid	1567
disable vrrp vrid	1568
enable vrrp vrid	1569
show vrrp	1570

show vrrp virtual-router	1573
show vrrp virtual-router	1575
show vrrp vlan.....	1577
Chapter 27: MPLS Commands—BlackDiamond 10808 and 12000 series Switches Only	1579
clear counters mpls	1580
clear counters mpls ldp	1581
clear counters mpls rsvp-te	1582
clear counters vpls	1583
clear fdb vpls.....	1584
configure fdb vpls agingtime	1585
configure iproute add	1586
configure iproute add default	1588
configure iproute delete.....	1590
configure iproute delete default.....	1591
configure mpls add vlan.....	1592
configure mpls delete vlan	1593
configure mpls exp examination.....	1594
configure mpls exp replacement	1595
configure mpls ldp advertise	1596
configure mpls ldp loop-detection.....	1598
configure mpls ldp timers	1599
configure mpls lsr-id	1601
configure mpls rsvp-te bandwidth committed-rate	1602
configure mpls rsvp-te lsp add path	1604
configure mpls rsvp-te lsp delete path.....	1606
configure mpls rsvp-te lsp path use profile.....	1607
configure mpls rsvp-te lsp transport	1608
configure mpls rsvp-te metric	1610
configure mpls rsvp-te path add ero	1611
configure mpls rsvp-te path delete ero.....	1613
configure mpls rsvp-te profile	1614
configure mpls rsvp-te timers lsp rapid-retry	1617
configure mpls rsvp-te timers lsp standard-retry	1619
configure mpls rsvp-te timers session	1621
configure vpls	1624
configure vpls add peer.....	1626
configure vpls add service	1627
configure vpls delete peer	1629
configure vpls delete service	1630
configure vpls peer	1631
configure vpls peer mpls lsp.....	1632
create mpls rsvp-te lsp	1633
create mpls rsvp-te path	1634
create mpls rsvp-te profile.....	1635
create vpls fec-id-type pseudo-wire	1636

delete mpls rsvp-te lsp	1637
delete mpls rsvp-te path	1638
delete mpls rsvp-te profile.....	1639
delete vpls.....	1640
disable bgp mpls-next-hop	1641
disable iproute mpls-next-hop	1642
disable mpls.....	1643
disable mpls exp examination.....	1644
disable mpls exp replacement	1645
disable mpls ldp	1646
disable mpls ldp bgp-routes	1647
disable mpls ldp loop-detection	1648
disable mpls php	1649
disable mpls protocol ldp.....	1650
disable mpls protocol rsvp-te.....	1651
disable mpls rsvp-te	1652
disable mpls rsvp-te bundle-message	1653
disable mpls rsvp-te lsp	1654
disable mpls rsvp-te summary-refresh	1655
disable mpls vlan	1656
disable ospf mpls-next-hop	1657
disable snmp traps mpls.....	1658
disable vpls	1659
disable vpls service	1660
enable bgp mpls-next-hop.....	1661
enable iproute mpls-next-hop	1662
enable mpls.....	1663
enable mpls exp examination	1664
enable mpls exp replacement	1665
enable mpls ldp	1666
enable mpls ldp bgp-routes.....	1667
enable mpls ldp loop-detection.....	1668
enable mpls php	1669
enable mpls protocol ldp.....	1670
enable mpls protocol rsvp-te	1671
enable mpls rsvp-te	1672
enable mpls rsvp-te bundle-message.....	1673
enable mpls rsvp-te lsp.....	1674
enable mpls rsvp-te summary-refresh	1675
enable mpls vlan.....	1676
enable ospf mpls-next-hop	1677
enable snmp traps mpls.....	1678
enable vpls.....	1679
enable vpls service	1680
ping mpls lsp.....	1681
restart process mpls	1684

show bandwidth pool	1685
show fdb	1687
show mpls	1689
show mpls exp examination	1691
show mpls exp replacement	1692
show mpls interface	1693
show mpls label	1695
show mpls ldp	1698
show mpls ldp interface	1700
show mpls ldp lsp	1702
show mpls ldp peer	1704
show mpls rsvp-te	1706
show mpls rsvp-te bandwidth	1708
show mpls rsvp-te interface	1710
show mpls rsvp-te lsp	1713
show mpls rsvp-te lsp (egress/transit)	1717
show mpls rsvp-te lsp (ingress)	1720
show mpls rsvp-te neighbor	1724
show mpls rsvp-te path	1726
show mpls rsvp-te profile	1727
show vpls	1729
show vpls peer	1731
traceroute mpls lsp	1733
unconfigure mpls	1735
unconfigure mpls exp examination	1736
unconfigure mpls exp replacement	1737
unconfigure mpls vlan	1738
unconfigure vpls dot1q ethertype	1739
Chapter 28: IP Unicast Commands	1741
clear iparp	1743
configure bootprelay add	1744
configure bootprelay delete	1745
configure bootprelay dhcp-agent information check	1746
configure bootprelay dhcp-agent information option	1747
configure bootprelay dhcp-agent information policy	1748
configure iparp add	1749
configure iparp add proxy	1750
configure iparp delete	1752
configure iparp delete proxy	1753
configure iparp max_entries	1754
configure iparp max_pending_entries	1755
configure iparp max_proxy_entries	1756
configure iparp timeout	1757
configure ipforwarding originated-packets	1758
configure iproute add	1759

configure iproute add blackhole.....	1760
configure iproute add blackhole default.....	1761
configure iproute add default	1762
configure iproute delete	1763
configure iproute delete blackhole	1764
configure iproute delete blackhole default	1765
configure iproute delete default	1766
configure iproute priority.....	1767
configure iproute sharing max-gateways.....	1769
configure irdp	1770
configure vlan add secondary-ipaddress.....	1771
configure vlan delete secondary-ipaddress	1772
configure vlan udp-profile	1773
disable bootp vlan	1775
disable bootprelay	1776
disable icmp address-mask	1777
disable icmp parameter-problem.....	1778
disable icmp port-unreachables	1779
disable icmp redirects	1780
disable icmp time-exceeded.....	1781
disable icmp timestamp.....	1782
disable icmp unreachable.....	1783
disable icmp useredirects	1784
disable iparp checking.....	1785
disable iparp refresh.....	1786
disable ipforwarding	1787
disable iproute compression	1788
disable ip-option loose-source-route	1789
disable ip-option record-route.....	1790
disable ip-option record-timestamp.....	1791
disable ip-option strict-source-route	1792
disable ip-option router-alert	1793
disable irdp	1794
disable udp-echo-server	1795
enable bootp vlan	1796
enable bootprelay.....	1797
enable icmp address-mask.....	1798
enable icmp parameter-problem	1799
enable icmp port-unreachables.....	1800
enable icmp redirects	1801
enable icmp time-exceeded.....	1802
enable icmp timestamp	1803
enable icmp unreachable	1804
enable icmp useredirects	1805
enable iparp checking	1806
enable iparp refresh	1807

enable ipforwarding.....	1808
enable ip-option loose-source-route.....	1810
enable ip-option record-route	1811
enable ip-option record-timestamp.....	1812
enable ip-option strict-source-route.....	1813
enable ip-option router-alert.....	1814
enable iproute compression.....	1815
enable iproute sharing	1816
enable irdp.....	1817
enable udp-echo-server.....	1818
rtlookup	1819
show bootprelay	1820
show iparp.....	1821
show iparp proxy	1824
show ipconfig	1825
show iproute.....	1826
show iproute origin	1828
show ipstats	1829
show udp-profile	1832
unconfigure bootprelay dhcp-agent information check	1833
unconfigure bootprelay dhcp-agent information option.....	1834
unconfigure bootprelay dhcp-agent information policy	1835
unconfigure icmp.....	1836
unconfigure iparp.....	1837
unconfigure irdp	1838
unconfigure udp-profile	1839
Chapter 29: IPv6 Unicast Commands.....	1841
clear neighbor-discovery cache	1842
configure iproute add	1843
configure iproute add blackhole.....	1844
configure iproute add blackhole ipv6 default	1845
configure iproute add default	1846
configure iproute delete.....	1847
configure iproute delete blackhole	1848
configure iproute delete blackhole ipv6 default	1849
configure iproute delete default.....	1850
configure iproute ipv6 priority	1851
configure neighbor-discovery cache add	1853
configure neighbor-discovery cache delete.....	1854
configure vlan router-discovery add prefix.....	1855
configure vlan router-discovery delete prefix.....	1856
configure vlan router-discovery default-lifetime	1857
configure vlan router-discovery hop-limit	1858
configure vlan router-discovery link-mtu	1859
configure vlan router-discovery managed-config-flag	1860

configure vlan router-discovery max-interval	1861
configure vlan router-discovery min-interval	1862
configure vlan router-discovery other-config-flag	1863
configure vlan router-discovery reachable-time	1864
configure vlan router-discovery retransmit-time	1865
configure vlan router-discovery set prefix	1866
configure tunnel ipaddress	1868
create tunnel 6to4	1869
create tunnel ipv6-in-ipv4	1870
delete tunnel	1871
disable ipforwarding ipv6	1872
disable iproute ipv6 compression	1873
disable router-discovery	1874
enable ipforwarding ipv6	1875
enable iproute ipv6 compression	1876
enable router-discovery	1877
rtlookup	1878
show ipconfig ipv6	1879
show iproute ipv6	1880
show iproute ipv6 origin	1881
show ipstats ipv6	1882
show neighbor-discovery cache ipv6	1883
show router-discovery	1885
show tunnel	1886
unconfigure tunnel	1887
Chapter 30: RIP Commands.....	1889
clear rip counters	1890
configure rip add vlan	1891
configure rip delete vlan	1892
configure rip garbagetime	1893
configure rip import-policy	1894
configure rip routetimeout	1895
configure rip vlan rxmode	1896
configure rip vlan txmode	1897
configure rip updatetime	1898
configure rip vlan cost	1899
configure rip vlan route-policy	1900
configure rip vlan trusted-gateway	1901
disable rip	1902
disable rip aggregation	1903
disable rip export	1904
disable rip originate-default	1905
disable rip poisonreverse	1906
disable rip splithorizon	1907
disable rip triggerupdates	1908

disable rip use-ip-router-alert	1909
enable rip	1910
enable rip aggregation	1911
enable rip export	1912
enable rip originate-default cost	1914
enable rip poisonreverse	1915
enable rip splithorizon	1916
enable rip triggerupdates	1917
enable rip use-ip-router-alert	1918
show rip	1919
show rip interface	1920
show rip interface vlan	1922
show rip memory	1923
show rip routes	1924
unconfigure rip	1925
Chapter 31: RIPng Commands	1927
clear ripng counters	1928
configure ripng add	1929
configure ripng delete	1930
configure ripng garbagetime	1931
configure ripng import-policy	1932
configure ripng routetimeout	1933
configure ripng updatetime	1934
configure ripng cost	1935
configure ripng route-policy	1936
configure ripng trusted-gateway	1938
disable ripng	1940
disable ripng export	1941
disable ripng originate-default	1942
disable ripng poisonreverse	1943
disable ripng splithorizon	1944
disable ripng triggerupdate	1945
enable ripng	1946
enable ripng export	1947
enable ripng originate-default	1949
enable ripng poisonreverse	1950
enable ripng splithorizon	1951
enable ripng triggerupdates	1952
show ripng	1953
show ripng interface	1954
show ripng routes	1956
unconfigure ripng	1958
Chapter 32: OSPF Commands	1959
clear ospf counters	1961

configure ospf cost	1962
configure ospf priority	1963
configure ospf authentication	1964
configure ospf add virtual-link	1966
configure ospf add vlan area	1967
configure ospf add vlan area link-type	1968
configure ospf area external-filter	1969
configure ospf area interarea-filter	1970
configure ospf area add range	1971
configure ospf area delete range	1972
configure ospf area normal	1973
configure ospf area nssa stub-default-cost	1974
configure ospf area stub stub-default-cost	1975
configure ospf area timer	1976
configure ospf ase-limit	1978
configure ospf ase-summary add	1979
configure ospf ase-summary delete	1980
configure ospf delete virtual-link	1981
configure ospf delete vlan	1982
configure ospf import-policy	1983
configure ospf lsa-batch-interval	1984
configure ospf metric-table	1985
configure ospf restart	1986
configure ospf restart grace-period	1987
configure ospf restart-helper	1988
configure ospf routerid	1990
configure ospf spf-hold-time	1991
configure ospf virtual-link timer	1992
configure ospf vlan area	1994
configure ospf vlan neighbor add	1995
configure ospf vlan neighbor delete	1996
configure ospf vlan timer	1997
create ospf area	1999
delete ospf area	2000
disable ospf	2001
disable ospf capability opaque-lsa	2002
disable ospf export	2003
disable ospf originate-default	2004
disable ospf restart-helper-lsa-check	2005
disable ospf use-ip-router-alert	2006
enable ospf	2007
enable ospf capability opaque-lsa	2008
enable ospf export	2009
enable ospf originate-default	2011
enable ospf restart-helper-lsa-check	2012
enable ospf use-ip-router-alert	2013

show ospf	2014
show ospf area	2015
show ospf area detail	2016
show ospf ase-summary	2017
show ospf interfaces detail	2018
show ospf interfaces	2019
show ospf lsdb	2020
show ospf memory	2022
show ospf neighbor	2023
show ospf virtual-link	2024
unconfigure ospf	2025
Chapter 33: OSPFv3 Commands	2027
clear ospfv3 counters	2028
configure ospfv3 add interface	2030
configure ospfv3 add interface all	2032
configure ospfv3 add virtual-link	2034
configure ospfv3 area add range	2035
configure ospfv3 area cost	2036
configure ospfv3 area delete range	2037
configure ospfv3 area external-filter	2038
configure ospfv3 area interarea-filter	2040
configure ospfv3 area normal	2042
configure ospfv3 area priority	2043
configure ospfv3 area stub	2044
configure ospfv3 area timer	2045
configure ospfv3 delete interface	2047
configure ospfv3 delete virtual-link	2048
configure ospfv3 import-policy	2049
configure ospfv3 interface area	2050
configure ospfv3 interface cost	2051
configure ospfv3 interface priority	2052
configure ospfv3 interface timer	2053
configure ospfv3 metric-table	2055
configure ospfv3 routerid	2056
configure ospfv3 spf-hold-time	2057
configure ospfv3 virtual-link timer	2058
create ospfv3 area	2060
delete ospfv3 area	2061
disable ospfv3	2062
disable ospfv3 export	2063
enable ospfv3	2064
enable ospfv3 export	2065
show ospfv3	2067
show ospfv3 area	2069
show ospfv3 interfaces	2071

show ospfv3 lsdb	2073
show ospfv3 lsdb stats.....	2075
show ospfv3 memory	2077
show ospfv3 neighbor	2078
show ospfv3 virtual-link	2079
unconfigure ospfv3.....	2080
Chapter 34: BGP Commands.....	2083
clear bgp neighbor counters	2085
clear bgp flap-statistics	2086
configure bgp add aggregate-address	2088
configure bgp add confederation-peer sub-AS-number	2090
configure bgp add network	2091
configure bgp AS-number	2092
configure bgp cluster-id	2093
configure bgp confederation-id	2094
configure bgp delete aggregate-address	2095
configure bgp delete confederation-peer sub-AS-number	2096
configure bgp delete network.....	2097
configure bgp export shutdown-priority	2098
configure bgp import-policy.....	2100
configure bgp local-preference	2101
configure bgp med	2102
configure bgp neighbor dampening	2103
configure bgp neighbor description	2105
configure bgp neighbor maximum-prefix	2106
configure bgp neighbor next-hop-self	2108
configure bgp neighbor no-dampening.....	2109
configure bgp neighbor password	2110
configure bgp neighbor peer-group.....	2112
configure bgp neighbor route-policy	2114
configure bgp neighbor route-reflector-client.....	2115
configure bgp neighbor send-community	2116
configure bgp neighbor shutdown-priority	2117
configure bgp neighbor soft-reset.....	2118
configure bgp neighbor source-interface	2120
configure bgp neighbor timer.....	2121
configure bgp neighbor weight.....	2122
configure bgp peer-group dampening	2123
configure bgp peer-group maximum-prefix	2125
configure bgp peer-group next-hop-self.....	2127
configure bgp peer-group no-dampening.....	2128
configure bgp peer-group route-reflector-client.....	2129
configure bgp peer-group send-community	2130
configure bgp peer-group password	2131
configure bgp peer-group remote-AS-number	2132

configure bgp peer-group route-policy	2133
configure bgp peer-group soft-reset.....	2134
configure bgp peer-group source-interface	2136
configure bgp peer-group timer.....	2138
configure bgp peer-group weight.....	2140
configure bgp restart	2141
configure bgp restart address-family.....	2142
configure bgp restart restart-time.....	2143
configure bgp restart stale-route-time.....	2144
configure bgp restart update-delay	2145
configure bgp routerid	2146
configure bgp soft-reconfiguration.....	2147
create bgp neighbor peer-group	2148
create bgp neighbor remote-AS-number.....	2149
create bgp peer-group.....	2150
delete bgp neighbor.....	2151
delete bgp peer-group.....	2152
disable bgp	2153
disable bgp aggregation	2154
disable bgp always-compare-med.....	2155
disable bgp community format	2156
disable bgp export.....	2157
disable bgp neighbor	2159
disable bgp neighbor capability	2160
disable bgp neighbor remove-private-AS-numbers.....	2162
disable bgp neighbor soft-in-reset	2163
disable bgp neighbor use-ip-router-alert	2164
disable bgp peer-group	2165
disable bgp peer-group capability	2166
disable bgp peer-group remove-private-AS-numbers.....	2167
disable bgp peer-group soft-in-reset	2168
disable bgp peer-group use-ip-router-alert.....	2169
enable bgp	2170
enable bgp aggregation.....	2171
enable bgp always-compare-med	2172
enable bgp community format	2173
enable bgp export	2174
enable bgp neighbor	2176
enable bgp neighbor capability	2177
enable bgp neighbor remove-private-AS-numbers	2179
enable bgp neighbor soft-in-reset.....	2180
enable bgp neighbor use-ip-router-alert	2181
enable bgp peer-group	2182
enable bgp peer-group capability	2183
enable bgp peer-group remove-private-AS-numbers	2184
enable bgp peer-group soft-in-reset.....	2185

enable bgp peer-group use-ip-router-alert	2186
show bgp	2187
show bgp neighbor	2188
show bgp neighbor flap-statistics	2189
show bgp neighbor routes	2191
show bgp peer-group	2193
show bgp routes	2194
show bgp routes summary	2196
show bgp memory	2197
Chapter 35: IP Multicast Commands	2199
clear igmp group	2202
clear igmp snooping	2203
clear pim cache	2204
configure igmp	2205
configure igmp snooping flood-list	2207
configure igmp snooping leave-timeout	2209
configure igmp snooping timer	2210
configure igmp snooping vlan ports add static group	2211
configure igmp snooping vlan ports delete static group	2213
configure igmp snooping vlan ports add static router	2214
configure igmp snooping vlan ports delete static router	2215
configure igmp snooping vlan ports filter	2216
configure ipmroute add	2218
configure ipmroute delete	2220
configure mvr add receiver	2221
configure mvr add vlan	2222
configure mvr delete receiver	2223
configure mvr delete vlan	2224
configure mvr mvr-address	2225
configure mvr static group	2227
configure pim add vlan	2229
configure pim cbsr	2231
configure pim crp static	2232
configure pim crp timer	2234
configure pim crp vlan	2235
configure pim delete vlan	2237
configure pim register-rate-limit-interval	2238
configure pim register-suppress-interval register-probe-interval	2239
configure pim register-checksum-to	2240
configure pim spt-threshold	2241
configure pim ssm range	2242
configure pim timer vlan	2244
configure pim vlan trusted-gateway	2245
disable igmp	2246
disable igmp snooping	2247

disable ipmcforwarding.....	2249
disable mvr	2250
disable pim	2251
disable pim ssm vlan.....	2252
enable igmp	2253
enable igmp snooping.....	2254
enable igmp snooping with-proxy	2256
enable ipmcforwarding	2257
enable mvr	2258
enable pim	2259
enable pim ssm vlan.....	2260
show igmp.....	2261
show igmp group.....	2263
show igmp snooping	2265
show igmp snooping vlan filter	2267
show igmp snooping vlan static	2268
show ipmroute	2269
show L2stats	2271
show mvr.....	2272
show mvr cache	2273
show pim	2274
show pim cache	2277
unconfigure igmp.....	2279
unconfigure pim.....	2280
unconfigure pim ssm range	2281
Chapter 36: IPv6 Multicast Commands	2283
clear mld group.....	2284
clear mld snooping.....	2285
configure mld	2286
configure mld snooping vlan ports add static group	2287
configure mld snooping vlan ports delete static group	2288
configure mld snooping vlan ports add static router	2289
configure mld snooping vlan ports delete static router.....	2290
configure mld snooping flood-list.....	2291
configure mld snooping leave-timeout	2293
configure mld snooping timer	2294
disable mld	2296
disable mld snooping.....	2297
enable mld	2298
enable mld snooping	2299
enable mld snooping with-proxy.....	2301
show mld	2302
show mld group	2304
show mld snooping.....	2305
show mld snooping vlan static	2307

unconfigure mld.....	2308
Chapter 37: MSDP Commands.....	2309
clear msdp counters	2310
clear msdp sa-cache.....	2312
configure msdp default-peer.....	2313
configure msdp max-rejected-cache	2315
configure msdp mesh-group	2316
configure msdp no-default-peer	2318
configure msdp originator-id.....	2319
configure msdp password.....	2320
configure msdp peer description.....	2322
configure msdp sa-cache-server	2323
configure msdp sa-filter	2325
configure msdp sa-limit	2327
configure msdp source-interface.....	2328
configure msdp timer	2330
configure msdp ttl-threshold	2331
configure pim border	2332
create msdp mesh-group.....	2333
create msdp peer	2334
delete msdp mesh-group.....	2335
delete msdp peer	2336
disable msdp	2337
disable msdp data-encapsulation.....	2338
disable msdp export local-sa	2339
disable msdp peer.....	2340
disable msdp process-sa-request	2341
enable msdp.....	2343
enable msdp data-encapsulation.....	2344
enable msdp export local-sa	2345
enable msdp peer	2347
enable msdp process-sa-request	2348
show msdp	2350
show msdp memory.....	2351
show msdp mesh-group	2353
show msdp peer	2355
show msdp sa-cache	2357
unconfigure msdp sa-cache-server	2359
unconfigure pim border	2360
Appendix A: Configuration and Image Commands	2361
configure firmware	2362
download bootrom.....	2365
download image.....	2368
install bootrom.....	2380
install firmware.....	2383

install image.....	2387
load script	2391
run update	2393
save configuration	2394
show configuration	2397
show memorycard	2400
synchronize	2401
synchronize stacking	2404
unconfigure switch	2406
uninstall image	2408
upload configuration.....	2410
use configuration	2415
use image	2417
Appendix B: Troubleshooting Commands	2419
clear elrp counters	2421
configure debug core-dumps	2422
configure elrp-client one-shot.....	2424
configure elrp-client periodic.....	2426
configure forwarding hash-algorithm	2428
disable elrp-client	2430
disable log debug-mode.....	2431
eject memorycard.....	2432
enable elrp-client.....	2433
enable log debug-mode.....	2434
nslookup	2435
run diagnostics	2437
run elrp	2442
save debug tracefiles memorycard	2444
show debug	2445
show diagnostics.....	2447
show elrp	2451
show forwarding configuration	2453
show tech.....	2454
top	2456
unconfigure elrp-client	2457
upload debug.....	2458
Appendix C: CNA Agent Commands.....	2461
clear cna-testplug counters	2462
configure cna-testplug scheduler ipaddress.....	2463
configure cna-testplug vlan	2465
disable cna-testplug	2466
enable cna-testplug.....	2467
show cna-testplug	2468
Index of Commands	2471



Preface

This preface provides an overview of this guide, describes guide conventions, and lists other publications that may be useful.

Introduction

This guide provides the complete syntax for all the commands available in the currently supported versions of the ExtremeXOS™ software running on modular switches from Extreme Networks®.

This guide is intended for use as a reference by network administrators who are responsible for installing and setting up network equipment. It assumes knowledge of Extreme Networks switch configuration. For conceptual information and guidance on configuring Extreme Networks switches, see the *ExtremeXOS Concepts Guide* for your version of the ExtremeXOS software.

Terminology

When features, functionality, or operation is specific to a switch family, the family name is used. Explanations about features and operations that are the same across all product families simply refer to the product as the “switch.”

Conventions

This sections discusses the conventions used in this guide.

Platform-dependent Conventions

Each command has a separate entry for platform availability, which addresses which platforms support the entire feature.

In many cases, although the command is available on all platforms, each platform uses specific keywords. These keywords specific to each platform are shown in the Syntax Description and discussed in the Usage Guidelines.

Text Conventions

Table 1 and Table 2 list conventions that are used throughout this guide.

Table 1: Notice icons

Icon	Notice Type	Alerts you to...
	Note	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.

Table 2: Text conventions

Convention	Description
Screen displays	This typeface indicates command syntax, or represents information as it appears on the screen.
The words “enter” and “type”	When you see the word “enter” in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says “type.”
[Key] names	Key names are written with brackets, such as [Return] or [Esc]. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press [Ctrl]+[Alt]+[Del].
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text.

Command Titles

For clarity and brevity, the command titles omit variables, values, and optional arguments. The complete command syntax is displayed directly below the command titles.

Related Publications

The publications related to this one are:

- ExtremeXOS release notes
- *ExtremeXOS Concepts Guide*
- *Extreme Networks Consolidated XOS Hardware Installation Guide*

Documentation for Extreme Networks products is available on the World Wide Web at the following location:

<http://www.extremenetworks.com/>

Using ExtremeXOS Publications Online

You can access ExtremeXOS publications by downloading them from the Extreme Networks World Wide Web location or from your ExtremeWare product CD. Publications are provided in Adobe® Portable Document Format (PDF). Displaying or printing PDF files requires that your computer be equipped with Adobe Reader® software, which is available free of charge from Adobe Systems Incorporated.

The user guide PDF file provides links that connect you directly to relevant command information in the command reference guide PDF file. This quick-referencing capability enables you to easily find detailed information in the command reference guide for any command mentioned in the user guide.

To ensure that the quick-referencing feature functions properly:

- 1 Download both the user guide PDF file and the command reference guide PDF file to the *same* destination directory on your computer.
- 2 You may open one or both PDF files. To enable cross-referenced linking between the user guide and command reference guide; however, it is recommended that for ease of use, you keep both files open concurrently on your computer desktop.



NOTE

If you activate a cross-referencing link from the concepts guide PDF file to the command reference PDF file when the command reference PDF file is closed (that is, not currently open on your computer desktop), the system will close the user guide PDF file and open the command reference PDF file. To keep both PDF files open when you activate a cross-reference link, open both PDF files before using the link.

All of these documents are available in Adobe PDF format. You must have Acrobat Reader 5.0 or later to properly open the documents. You must have Acrobat Reader 6.0 or later to use the cross-reference linking feature from the *ExtremeXOS Concepts Guide* to the *ExtremeXOS Command Reference Guide*.

Introduction

This guide provides details of the command syntax for all ExtremeXOS™ commands as of ExtremeXOS version 12.0.

This guide does not provide feature descriptions, explanations of the technologies, or configuration examples. For information about the various features and technologies supported by Extreme Networks switches, see the installation and user guides for your product. This guide does not replace the installation and user guides; this guide supplements the installation and user guides.

This chapter discusses the following topics:

- [Audience on page 43](#)
- [Structure of this Guide on page 44](#)
- [Platforms and Required Software Versions on page 44](#)
- [Software Licenses and Feature Packs on page 46](#)
- [Understanding the Command Syntax on page 50](#)
- [Port Numbering on page 52](#)
- [Line-Editing Keys on page 53](#)
- [Command History on page 53](#)

Audience

This guide is intended for use by network administrators who are responsible for installing and setting up network equipment. It assumes a basic working knowledge of the following:

- Local area networks (LANs)
- Ethernet concepts
- Ethernet switching and bridging concepts
- Routing concepts
- Internet Protocol (IP) concepts
- Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) concepts
- Border Gateway Protocol (BGP-4) concepts
- IP Multicast concepts
- Protocol Independent Multicast (PIM) concepts
- Simple Network Management Protocol (SNMP)

This guide also assumes that you have read the Installation and User Guide for your product.

Structure of this Guide

This guide documents each ExtremeXOS command. Related commands are grouped together and organized into chapters based on their most common usage. The chapters reflect the organization of the *ExtremeXOS Concepts Guide*. If a specific command is relevant to a wide variety of functions and could be included in a number of different chapters, we have attempted to place the command in the most logical chapter. Within each chapter, commands appear in alphabetical order. You can use the Index of Commands to locate specific commands if they do not appear where you expect to find them.

For each command, the following information is provided:

- **Command Syntax**—The actual syntax of the command. The syntax conventions (the use of braces or curly brackets, for example) are defined in the section “[Understanding the Command Syntax](#)” on page 50.
- **Description**—A brief (one sentence) summary of what the command does.
- **Syntax Description**—The definition of any keywords and options used in the command.
- **Default**—The defaults, if any, for this command. The default can be the default action of the command if optional arguments are not provided, or it can be the default state of the switch (such as for an enable/disable command).
- **Usage Guidelines**—Information to help you use the command. This may include prerequisites, prohibitions, and related commands, as well as other information.
- **Example**—Examples of the command usage, including output, if relevant.
- **History**—The version of ExtremeXOS in which the command was introduced, and version(s) where it was modified, if appropriate.

Platforms and Required Software Versions

ExtremeXOS is the full-featured software operating system that is designed to run on the Extreme Networks devices.

ExtremeXOS supports the following platforms:

- BlackDiamond® 10808 switch—ExtremeXOS 10.1 and higher
 - MSM 1
 - MSM 1XL
- BlackDiamond 12800 series switches—ExtremeXOS 11.4.1 and higher
 - MSM 5
 - MSM 5R



CAUTION

You cannot mix BlackDiamond 12800 R-series module with non-R-series modules in the BlackDiamond 12800 series switches. If you attempt to mix these, the system returns an error message.

- BlackDiamond 8810 switch (formerly known as Aspen)—ExtremeXOS 11.1 and higher
 - G48T module—ExtremeXOS 11.1 and higher
 - G48P module—ExtremeXOS 11.1 and higher

- G24X module—ExtremeXOS 11.1 and higher
- 10G4X module—ExtremeXOS 11.1 and higher

These modules are hereinafter referred to as BlackDiamond 8800 original modules.

- BlackDiamond 8806 switch—ExtremeXOS 11.3.1. and higher
The BlackDiamond 8806 and 8810 switch and all modules that function with both chassis are hereinafter referred as the BlackDiamond 8800 series switch.
- MSMs for the BlackDiamond 8800 series switch
 - MSM-G8X—ExtremeXOS 11.1 and higher
 - MSM-48—ExtremeXOS 11.6 and higher
- BlackDiamond a-series modules (function with BlackDiamond 8806 and 8810 chassis)
 - G48Ta module—ExtremeXOS 11.5 and higher
 - G48Xa module—ExtremeXOS 11.5 and higher
- BlackDiamond e-series modules (function with BlackDiamond 8806 and 8810 chassis)—ExtremeXOS 11.5 and higher
 - G48Te module
 - G48Pe module
- Summit® X450 series switch
 - Summit X450-24t - ExtremeXOS 11.2 and higher
 - Summit X450-24x - ExtremeXOS 11.2 and higher
- Summit X450a series switches
 - Summit X450a-24t—ExtremeXOS 11.5 and higher
 - Summit X450a-48t—ExtremeXOS 11.5 and higher
 - Summit X450a-24t DC—ExtremeXOS 11.5 and higher
 - Summit X450a-24x—ExtremeXOS 11.6 and higher
 - Summit X450a-24x DC—ExtremeXOS 11.6 and higher
- Summit X450e series switches
 - Summit X450e-24p—ExtremeXOS 11.5 and higher
 - Summit X450e-48p—ExtremeXOS 11.6 and higher
- Summit X250e series switches
 - X250e-48t, 48 10/100M ports copper
 - X250e-48p, 48 10/100M ports copper with PoE capability
 - X250e-24t, 24 10/100M ports copper
 - X250e-24p, 24 10/100M ports copper with PoE capability
 - X250e-24x, 24 100M ports fiber (SFP)
 - X250e-24xTAA, 24 100M ports fiber (SFP) for government

The term “Summit family” refers to all Summit switches; “Summit X450” series refers to either the X450-24t or the X450-24x switch. The term *SummitStack* refers to any combination of Summit switches arranged in a stack and operating in stacking mode. SummitStacks are supported in ExtremeXOS version 12.0 and higher.

**NOTE**

See the Extreme Networks Consolidated XOS Hardware Installation Guide for complete information about Extreme Networks devices running ExtremeXOS software.

Software Licenses and Feature Packs

**NOTE**

Information on required software licensing and feature packs is in the Platform Availability section for each command. If there is no licensing information for the command, you do not need a license.

This section discusses the following topics:

- [Software Licensing Overview on page 46](#)
- [Software License Levels on page 47](#)
- [Security Licensing on page 49](#)
- [Feature Packs Overview on page 49](#)
- [Displaying Software Licenses and Feature Packs on page 49](#)

**NOTE**

Refer to [Chapter 2, “Commands for Accessing the Switch”](#) for specific commands to enable and display the licenses and feature packs.

You can enable the software or feature pack license key one switch at a time, or you can download a license file onto the switch and then enable the licenses through file upload. (The license file has the extension <.xlic>.) You download the file using TFTP into the switch. The file can contain licenses for some or all of the Extreme Networks switches that the customer owns. During upload, only those license keys destined for the specific switch are used to attempt enabling the licenses. The license file is a text file that has the switch serial number, software license type, and license key; it is removed from the switch after the licenses are enabled.

Software Licensing Overview

As all switches ship with Edge capabilities, three levels of software licensing apply to ExtremeXOS version 12.0: the Edge, the Advanced Edge, the Core, and the Advanced Core license (refer to [Table 3](#)). Additionally, the U.S. government requires a security license to enable certain features.

Table 3: Licensing requirements for ExtremeXOS version 12.0

License level	BlackDiamond 10808 switch	BlackDiamond 12800 series switches	Summit X450e switches	Summit X450 and X450a switches	BlackDiamond 8800 series switch
Edge	Ships (as part of Core license)		Ships	Ships	Ships

Table 3: Licensing requirements for ExtremeXOS version 12.0 (Continued)

License level	BlackDiamond 10808 switch	BlackDiamond 12800 series switches	Summit X450e switches	Summit X450 and X450a switches	BlackDiamond 8800 series switch
Advanced Edge	Ships (as part of Core license)		Requires license	Ships	Ships
Core	Ships	- MSM-5 ships - MSM-5R ships	N/A	Requires license	Requires license
Advanced Core	- MSM 1—N/A - MSM1XL—Ships	N/A	N/A	N/A	N/A

With ExtremeXOS software version 11.2, BGP functionality moved from the Advanced Core license level to the Core license. You have BGP functionality with a Core license.

When you are working with modular switches, the license belongs with the switch chassis, not with the particular MSM module.

The licensing levels required for each feature, if any, are outlined in the discussion of each feature. If licensing is not mentioned, the full functionality of that feature is available on every switch.

If you attempt to execute a command and you do not either have the required license or have reached the limits defined by the current license level, the system returns one of the following messages:

```
Error: This command cannot be executed at the current license level.
Error: You have reached the maximum limit for this feature at this license level.
```

You can obtain a trial license, which allows you to use the license for either 30, 60, or 90 days. After you enable the trial license, the switch behaves as if all software license levels and feature packs were enabled. Once you have a trial license for any one of these periods, you cannot extend the time of the trial. And, trial licenses can be applied only once for each software version; if you upgrade to a different software version, you can reapply your trial license.

When operating in a SummitStack, each switch must be separately licensed. If the license levels on all switches in the stack are not all the same, the user can force the stack to operate at the minimum of these levels so that all nodes will operate.

Software License Levels

This section discusses the software licensing capabilities of the various switches.

Upgrading License on the Summit X450e Series Switch Only

The Summit X450e series switch ships with an Edge license. With ExtremeXOS 11.5, you can obtain the Advanced Edge license for these switches. The Core and Advanced Core licenses are not available for the Summit X450e series switch.

**NOTE**

Refer to the specific chapter of the ExtremeXOS Concepts Guide to determine if the Advanced Edge license is required for some functionality. If not noted, all functionality is available, and a license is not required. Information on required licensing is in the Platform Availability section in this guide for each command. If there is no licensing information for the command, you do not need a license.

Upgrading License on the BlackDiamond 8800 Series Switch and Summit X450 and X450a Series Switches Only

The BlackDiamond 8800 series switch and the Summit X450 and X450a series switches ship with an Advanced Edge license. With ExtremeXOS 11.5, you can obtain the Core license for these switches. The Core license provides additional functionality for some features, as well as Border Gateway Protocol (BGP) functionality, on these switches.

The Advanced Core license is not available for the BlackDiamond 8800 series switch or the Summit X450 or X450a series switches.

**NOTE**

Refer to the specific chapter of the ExtremeXOS Concepts Guide to determine if the Core license is required for some functionality. If not noted, all functionality is available, and a license is not required. Information on required licensing is in the Platform Availability section in this guide for each command. If there is no licensing information for the command, you do not need a license.

The BlackDiamond 10808 Switch Only

The licensing levels on the BlackDiamond 10808 switch depend on the MSM you have in your system:

- The MSM 1 ships with a Core license; you cannot upgrade to an Advanced Core license.
- The MSM 1XL ships with an Advanced Core license.

**NOTE**

Beginning with ExtremeXOS version 11.2, you can run BGP with a Core license on the BlackDiamond 10808 switch; that is both the MSM1 and the MSM1XL now run BGP.

The BlackDiamond 12800 Series Switches Only

Both the MSM-5 and the MSM-5R ship with a Core license, and the Advanced Core license is not available for this switch.

**NOTE**

Beginning with ExtremeXOS version 11.2, you can run BGP with a Core license.

Security Licensing

Certain additional ExtremeXOS features, such as the use of SSH2 encryption, may be under United States export restriction control. Once the export restrictions are met, you can obtain information on enabling these features at no charge from Extreme Networks.

Feature Packs Overview

In addition to licenses, some Extreme Networks features are available only by obtaining a separate feature pack. Like licenses, feature packs may be obtained from Extreme Networks at an additional cost. Feature packs must be purchased for each switch; they are tied to a switch and are not portable. Feature packs remain in place during software upgrades and or module replacements.

MPLS Feature Pack—BlackDiamond 10808 MSM-1XL and BlackDiamond 12800 R-series Switches Only

MPLS support, including VPLS L2 VPNs, may be obtained from Extreme Networks at an additional cost, as a feature pack. A separate feature pack must be purchased for each switch that will run MPLS and VPLS.

All the MPLS CLI commands are visible, even if the switch you are working on does not have an MPLS feature pack; you can configure MPLS and VPLS on a switch that does not have the feature pack. However, when you attempt to enable MPLS on a BlackDiamond 10808 MSM-1XL or BlackDiamond 12800 R-series switch without the feature pack, the following error message is displayed and MPLS is *not* enabled:

```
Error: MPLS feature not correctly licensed for this switch, CANNOT enable MPLS!
      Please contact an Extreme Networks sales representative or go
      to www.extremenetworks.com for additional information.
The show licenses command reveals that no feature packs have been licensed.
```

MEF Feature Pack—BlackDiamond 12800 R-series Switch Only

ExtremeXOS version 12.0 software introduces a feature pack for bandwidth mode hierarchical Quality of Service (QoS), which meets the standards of the Metro Ethernet Forum (MEF) 14. Strict priority mode hierarchical QoS, introduced in ExtremeXOS 11.4 does not require that you purchase a separate feature pack to implement.

If you do not have an enabled feature pack for bandwidth mode hierarchical QoS, you do not see any of the associated CLI commands.

Displaying Software Licenses and Feature Packs

You can display information on the type of license and feature pack on your Extreme Networks device by using the `show licenses` command:

```
BD-10808.4 # show licenses
Enabled License Level:
    Core
Enabled Feature Packs:
    MPLS
```

Understanding the Command Syntax

This section covers the following topics:

- [Access Levels on page 50](#)
- [Syntax Symbols on page 50](#)
- [Command Completion with Syntax Helper on page 51](#)
- [Names on page 51](#)
- [Command Shortcuts on page 51](#)

Access Levels

When entering a command at the prompt, ensure that you have the appropriate privilege level. Most configuration commands require you to have the administrator privilege level.

Syntax Symbols

You may see a variety of symbols shown as part of the command syntax. These symbols explain how to enter the command, and you do not type them as part of the command itself. [Table 4](#) summarizes command syntax symbols.

Table 4: Command syntax symbols

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax <pre>configure vlan <vlan_name> ipaddress <ip_address></pre> you must supply a VLAN name for <vlan_name> and an address for <ip_address> when entering the command. Do not type the angle brackets. You may not include spaces within angle brackets.
square brackets []	Enclose a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax <pre>use image [primary secondary]</pre> you must specify either the primary or secondary image when entering the command. Do not type the square brackets.
vertical bar	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax <pre>configure snmp community [readonly readwrite] <alphanumeric_string></pre> you must specify either the read or write community string in the command. Do not type the vertical bar.
braces { }	Enclose an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax <pre>reboot {time <month> <day> <year> <hour> <min> <sec>} {cancel} {msm <slot_id>} {slot <slot-number> node-address <node-address> stack-topology {as-standby} }</pre> you can specify either a particular date and time combination, or the keyword <code>cancel</code> to cancel a previously scheduled reboot. If you do not specify an argument, the command will prompt asking if you want to reboot the switch now. Do not type the braces.

Command Completion with Syntax Helper

The CLI has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press [Tab]. The syntax helper provides a list of options for the remainder of the command, and places the cursor at the end of the command you have entered so far, ready for the next option.

If the command is one where the next option is a named component, such as a VLAN, access profile, or route map, the syntax helper will also list any currently configured names that might be used as the next option. In situations where this list might be very long, the syntax helper will list only one line of names, followed by an ellipses to indicate that there are more names than can be displayed.

Some values (such as the <node-address> used in Summit stack) are lengthy, but limited in number. ExtremeXOS places these values into a “namespace.” This allows command completion on these values.

The syntax helper also provides assistance if you have entered an incorrect command.

Abbreviated Syntax

Abbreviated syntax is the shortest unambiguous allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command. If you do not enter enough letters to allow the switch to determine which command you mean, the syntax helper will provide a list of the options based on the portion of the command you have entered.



NOTE

When using abbreviated syntax, you must enter enough characters to make the command unambiguous and distinguishable to the switch.

Names

All named components within a category of the switch configuration, such as VLAN, must have a unique name. Names can be re-used across categories, however. Names must begin with an alphabetical character and cannot contain any spaces. The maximum length for a name is 32 characters. Names may contain alphanumeric characters and underscores (_) and cannot be keywords, such as vlan, stp, and so on.



NOTE

If you use the same name across categories (for example, STPD and EAPS names), you must specify the identifying keyword as well as the actual name.

Command Shortcuts

All named components within a category of the switch configuration must have a unique name. Components are named using the `create` command. When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, you must enter a unique VLAN name:

```
create vlan engineering
```

Once you have created the VLAN with a unique name, you can then eliminate the keyword `vlan` from all other commands that require the name to be entered (unless you used the same name for another category such as STPD or EAPS). For example, instead of entering the modular switch command:

```
configure vlan engineering delete port 1:3,4:6
```

you could enter the following shortcut:

```
configure engineering delete port 1:3,4:6
```

Port Numbering

Commands that require you to enter one or more port numbers use the parameter `<port_list>` in the syntax. The available variables differ on a stand-alone switch, modular switch, and SummitStack.



NOTE

The keyword `all` acts on all possible ports; it continues on all ports even if one port in the sequence fails.

Stand-alone Switch Numerical Ranges

Commands that require you to enter one or more port numbers on a stand-alone switch, such as a Summit family switch, use the parameter `<port_list>` in the syntax.

A `<port_list>` can be a single port. For example, port 5 would be:

```
5
```

A `<port_list>` can also be a range of numbers. Separate the port numbers by a dash to enter a range of contiguous numbers, and separate the numbers by a comma to enter a range of noncontiguous numbers:

- `x-y`—Specifies a contiguous series of ports on a stand-alone switch.
- `x,y`—Specifies a noncontiguous series of ports on a stand-alone switch.
- `x-y,a,d`—Specifies a contiguous series of ports and a noncontiguous series of ports on a stand-alone switch.

Modular Switch and SummitStack Numerical Ranges

Commands that require you to enter one or more port numbers on a modular switch, such as the BlackDiamond 10K switch, use the parameter `<port_list>` in the syntax. The syntax for the port and slot is:

```
port <slot_number>:<port_number>
```

A `<port_list>` can be one port on a particular slot. For example, port 1 on slot 3 would be:

```
port 3:1
```

A `<port_list>` can be a range of numbers. For example, ports 1 through 3 on slot 3 would be:

```
port 3:1-3:3
```


You can add additional slot and port numbers to the list, separated by a comma:

```
port 3:1,4:8,6:10
```

You can specify all ports on a particular slot, using the asterisk (*) wildcard. For example,

```
port 3:*
```

indicates all ports on slot 3.

You can specify a range of slots and ports. For example,

```
port 2:3-4:5
```

indicates slot 2, port 3 through slot 4, port 5.

Line-Editing Keys

Table 5 describes the line-editing keys available using the CLI.

Table 5: Line-editing keys

Key(s)	Description
Left arrow or [Ctrl] + B	Moves the cursor one character to the left.
Right arrow or [Ctrl] + F	Moves the cursor one character to the right.
[Ctrl] + H or Backspace	Deletes character to left of cursor and shifts remainder of line to left.
Delete or [Ctrl] + D	Deletes character under cursor and shifts remainder of line to left.
[Ctrl] + K	Deletes characters from under cursor to end of line.
Insert	Toggles on and off. When toggled on, inserts text and shifts previous text to right.
Left Arrow	Moves cursor to left.
Right Arrow	Moves cursor to right.
[Ctrl] + L	Clears screen and moves cursor to beginning of line.
[Ctrl] + P or Up Arrow	Displays previous command in command history buffer and places cursor at end of command.
[Ctrl] + N or Down Arrow	Displays next command in command history buffer and places cursor at end of command.
[Ctrl] + U	Clears all characters typed from cursor to beginning of line.
[Ctrl] + W	Deletes previous word.
[Ctrl] + C	Interrupts the current CLI command execution.

Command History

ExtremeXOS “remembers” all the commands you enter. You can display a list of these commands by using the following command:

```
history
```

If you use a command more than once, consecutively, the history will only list the first instance.

2

Commands for Accessing the Switch

This chapter describes commands used for:

- Accessing and configuring the switch including how to set up user accounts, passwords, date and time settings, and software licenses
- Managing passwords
- Configuring the Domain Name Service (DNS) client
- Checking basic switch connectivity
- Enabling and displaying licenses
- Returning the switch to safe defaults mode

ExtremeXOS supports the following two levels of management:

- User
- Administrator

A user-level account has viewing access to all manageable parameters, with the exception of:

- User account database
- SNMP community strings

A user-level account can change the password assigned to the account name and use the `ping` command to test device reachability.

An administrator-level account can view and change all switch parameters. It can also add and delete users and change the password associated with any account name. The administrator can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged on by way of the Telnet connection is notified that the session has been terminated.

The DNS client in ExtremeXOS augments certain ExtremeXOS commands to accept either IP addresses or host names. For example, DNS can be used during a Telnet session when you are accessing a device or when using the `ping` command to check the connectivity of a device.

The switch offers the following commands for checking basic connectivity:

- `ping`
- `traceroute`

The `ping` command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The `traceroute` command enables you to trace the routed path between the switch and a destination endstation.

This chapter describes commands for enabling and displaying software, security, and feature pack licenses.

clear account logout

```
clear account [all | <name>] logout
```

Description

Re-enables an account that has been locked out (disabled) for exceeding the permitted number failed login attempts, which was configured by using the `configure account password-policy logout-on-login-failures` command.

Syntax Description

all	Specifies all users.
name	Specifies an account name.

Default

N/A.

Usage Guidelines

This command applies to sessions at the console port of the switch as well as all other sessions. You can re-enable both user and administrative accounts, once they have been disabled for exceeding the 3 failed login attempts.



NOTE

The failsafe accounts are never locked out.

This command only clears the locked-out (or disabled) condition of the account. The action of locking out accounts following the failed login attempts *remains* until you turn it off by issuing the `configure account [all | <name>] password-policy logout-on-login failures off` command.

Example

The following command re-enables the account finance, which had been locked out (disabled) for exceeding 3 consecutive failed login attempts:

```
clear account finance logout
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

clear license-info

```
clear license-info
```

Description

This command, which should be used only in conjunction with a representative from Extreme Networks, clears the licensing information from the switch.

Syntax Description

This command has no variables or parameters.

Default

N/A.

Usage Guidelines



NOTE

Use this command only under the guidance of an Extreme Networks representative.

This command clears licensing information from the switch. When you issue this command, the system requests a confirmation. If you answer yes, the system sends a Warning message to the log.

Example

The following command removes licensing information from the switch:

```
clear license-info
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

clear session

```
clear session [<sessId> | all]
```

Description

Terminates a Telnet and/or SSH2 sessions from the switch.

Syntax Description

sessId	Specifies a session number from <code>show session</code> output to terminate.
all	Terminates all sessions.

Default

N/A.

Usage Guidelines

An administrator-level account can disconnect a management session that has been established by way of a Telnet connection. You can determine the session number of the session you want to terminate by using the `show session` command. The `show session` output displays information about current Telnet and/or SSH2 sessions including:

- The session number
- The login date and time
- The user name
- The type of Telnet session
- Authentication information

Depending on the software version running on your switch, additional session information may be displayed. The session number is the first number displayed in the `show session` output.

Example

The following command terminates session 4 from the system:

```
clear session 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure account

```
configure account [all | <name>]
```

Description

Configures a password for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password. Use the [create account](#) command to create a user account.

The system prompts you to specify a password after you enter this command. You must enter a password for this command; passwords cannot be null.



NOTE

Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the [create account](#) command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive; user names are not case-sensitive.



NOTE

If the account is configured to require a specific password format, the minimum is 8 characters. See [configure account password-policy char-validation](#) for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command defines a new password *green* for the account *marketing*:

```
configure account marketing
```

The switch responds with a password prompt:

```
password: green
```

Your keystrokes will not be echoed as you enter the new password. After you enter the password, the switch will then prompt you to reenter it.

```
Reenter password: green
```

Assuming you enter it successfully a second time, the password is now changed.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure account encrypted

```
configure account [all | <name>] encrypted <e-password>
```

Description

Encrypts the password that is entered in plain text for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
e-password	Enter in plain text the string you for an encrypted password. See "Usage Guidelines" for more information.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password. Use the `create account` account command to create a user account.

When you use this command, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should enter the password in plain text. The encrypted password is then used by the switch once it encrypts the plain text password. The encrypted command should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.



NOTE

Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive; user names are not case-sensitive.



NOTE

If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command encrypts the password *red* for the account *marketing*:

```
configure account marketing encrypted red
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure account password-policy char-validation

```
configure account [all | <name>] password-policy char-validation [none |
all-char-groups]
```

Description

Requires that the user include an upper-case letter, a lower-case letter, a digit, and a symbol in the password.

Syntax Description

all	Specifies all users (and future users).
name	Specifies an account name.
none	Resets password to accept all formats.
all-char-groups	Specifies that the password must contain at least two characters from each of the four groups. NOTE: The password minimum length will be 8 characters if you specify this option.

Default

N/A.

Usage Guidelines

This feature is disabled by default.

Once you issue this command, each password must include at least *two* characters of each of the following four types:

- Upper-case A-Z
- Lower-case a-z
- 0-9
- !, @, #, \$, %, ^, *, (,)

The minimum number of characters for these specifically formatted passwords is 8 characters and the maximum is 32 characters.

Use the `none` option to reset the password to accept all formats.

Example

The following command requires all users to use this specified format for all passwords:

```
configure account all password-policy char-validation all-char-groups
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy history

```
configure account [all | <name>] password-policy history [<num_passwords> | none]
```

Description

Configures the switch to verify the specified number of previous passwords for the account. The user is prevented from changing the password on a user or administrative account to any of these previously saved passwords.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
num_passwords	Specifies the number of previous passwords the system verifies for each account. The range is 1 to 10 passwords.
none	Resets the system to not remember any previous passwords.

Default

N/A.

Usage Guidelines

Use this command to instruct the system to verify new passwords against a list of all previously used passwords, once an account successfully changes a password. The limit is the number of previous passwords that the system checks against in the record to verify the new password.

If this parameter is configured, the system returns an error message if a user attempts to change the password to one that is saved by the system (up to the configured limit) for that account; this applies to both user and administrative accounts. This also applies to a configured password on the default admin account on the switch.

The limit of previous passwords that the system checks for previous use is configurable from 1 to 10. Using the `none` option disables previous password tracking and returns the system to the default state of no record of previous passwords.

Example

The following command instructs the system to verify that the new password has not been used as a password in the previous 5 passwords for the account engineering:

```
configure account engineering password-policy history 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy logout-on-login-failures

```
configure account [all | <name>] password-policy logout-on-login-failures
[on | off]
```

Description

Disables an account after the user has 3 consecutive failed login attempts.

Syntax Description

all	Specifies all users (and future users).
name	Specifies an account name.
on	Specifies an account name.
off	Resets the password to never lockout the user.

Default

N/A.

Usage Guidelines

If you are not working on SSH, you can configure the number of failed logins that trigger lockout, using the `configure cli max-failed-logins <num-of-logins>` command.

This command applies to sessions at the console port of the switch as well as all other sessions and to user-level and administrator-level accounts. This command locks out the user after 3 consecutive failed login attempts; the user's account must be specifically re-enabled by an administrator.

Using the `off` option resets the account to allow innumerable consecutive failed login attempts, which is the system default. The system default is that 3 failed consecutive login attempts terminate the particular session, but the user may launch another session; there is no lockout feature by default.



NOTE

The failsafe accounts are never locked out, no matter how many consecutive failed login attempts.

Example

The following command enables the account finance for lockout. After 3 consecutive failed login attempts, the account is subsequently locked out:

```
configure account finance password-policy logout-on-login-failures on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy max-age

```
configure account [all | <name>] password-policy max-age [<num_days> | none]
```

Description

Configures a time limit for the passwords for specified accounts. The passwords for the default admin account and the failsafe account do not age out.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
num_days	Specifies the length of time that a password can be used. The range is 1 to 365 days.
none	Resets the password to never expire.

Default

N/A.

Usage Guidelines

The passwords for the default admin account and the failsafe account never expire.

The time limit is specified in days, from 1 to 365 days. Existing sessions are not closed when the time limit expires; it will not open the next time the user attempts to log in.

When a user logs into an account with an expired password, the system first verifies that the entered password had been valid prior to expiring and then prompts the user to change the password.



NOTE

This is the sole time that a user with a user-level (opposed to an administrator-level) account can make any changes to the user-level account.

Using the `none` option prevents the password for the specified account from ever expiring (it resets the password to the system default of no time limit).

Example

The following command sets a 3-month time limit for the password for the account marketing:

```
configure account marketing password-policy max-age 90
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy min-length

```
configure account [all | <name>] password-policy min-length
[<num_characters> | none]
```

Description

Requires a minimum number of characters for passwords.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
num_characters	Specifies the minimum number of characters required for the password. The range is 1 to 32 characters. NOTE: If you configure the configure account password-policy char-validation parameter, the minimum length is 8 characters.
none	Resets password to accept a minimum of 0 characters. NOTE: If you configure the configure account encrypted parameter, the minimum length is 8 characters.

Default

N/A.

Usage Guidelines

Use this command to configure a minimum length restriction for all passwords for specified accounts. This command affects the minimum allowed length for the *next* password; the current password is unaffected.

The minimum password length is configurable from 1 to 32 characters. Using the `none` option disables the requirement of minimum password length and returns the system to the default state (password minimum is 0 by default).



NOTE

If the account is configured to require a specific password format, the minimum is 8 characters. See [configure account password-policy char-validation](#) for more information.

Example

The following command requires a minimum of 8 letters for the password for the account management:

```
configure account management password-policy min-length 8
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure banner

```
configure banner {acknowledge}
```

Description

Configures the banner string that is displayed at the beginning of each login prompt of each session.

Syntax Description

acknowledge	Specifies that the system return the user-defined message after the banner is displayed. The user must then press a key (any key) to accept before the login displays. Certain systems require this configuration (for example, the U.S. Department of Defense).

Default

N/A.

Usage Guidelines

Press [Return] at the beginning of a line to terminate the command and apply the banner. To clear the banner, press [Return] at the beginning of the first line. You can enter up to 24 rows of 79-column text that is displayed before the login prompt of each session. To disable the acknowledgement feature, use the `configure banner` command omitting the `acknowledge` parameter.



NOTE

The system does not wait for a keypress when you use SSH for access; this only applies to the serial console login sessions and telnet sessions.

Example

The following command adds a banner, *Welcome to the switch*, before the login prompt:

```
configure banner [Return]
Welcome to the switch
```

History

This command was first available in ExtremeXOS 10.1.

The `acknowledge` parameter was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure cli max-sessions

```
configure cli max-sessions <num-of-sessions>
```

Description

Limits number of simultaneous CLI sessions on the switch.

Syntax Description

num-of-sessions	Specifies the maximum number of concurrent sessions permitted. The range is 1 to 16.

Default

The default is eight sessions.

Usage Guidelines

The value must be greater than 0; the range is 1 to 16.

Example

The following command limits the number of simultaneous CLI sessions to ten:

```
configure cli max-sessions 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure cli max-failed-logins

```
configure cli max-failed-logins <num-of-logins>
```

Description

Establishes the maximum number of failed logins permitted before the session is terminated.

Syntax Description

num-of-logins	Specifies the maximum number of failed logins permitted; the range is 1 to 10.

Default

The default is three logins.

Usage Guidelines

The value must be greater than 0; the range is 1 to 10.

Example

The following command sets the maximum number of failed logins to five:

```
configure cli max-failed-logins 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure dns-client add

```
configure dns-client add [domain-suffix <domain_name> | name-server
<ip_address> {vr <vr_name>}]
```

Description

Adds a domain suffix to the domain suffix list or a name server to the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies adding a domain suffix.
domain_name	Specifies a domain name.
name-server	Specifies adding a name server.
ip_address	Specifies an IP address for the name server.
vr	Specifies use of a virtual router.
	NOTE: User-created VRs are not available on the BlackDiamond 8800 series switch (formerly known as Aspen) or the Summit family of switches.
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

The domain suffix list can include up to six items. If the use of all previous names fails to resolve a name, the most recently added entry on the domain suffix list will be the last name used during name resolution. This command will not overwrite any exiting entries. If a null string is used as the last suffix in the list, and all other lookups fail, the name resolver will attempt to look up the name with no suffix.

Up to eight DNS name servers can be configured. The default value for the virtual router used by the DNS client option is VR-Default.

Examples

The following command configures a domain name and adds it to the domain suffix list:

```
configure dns-client add domain-suffix xyz_inc.com
```

The following command specifies that the switch use the DNS server 10.1.2.1:

```
configure dns-client add name-server 10.1.2.1
```

The following command on the BlackDiamond 10808 switch specifies that the switch use the virtual router Management:

```
configure dns-client add name-server 10.1.2.1 vr "VR-Management"
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure dns-client default-domain

```
configure dns-client default-domain <domain_name>
```

Description

Configures the domain that the DNS client uses if a fully qualified domain name is not entered.

Syntax Description

domain_name	Specifies a default domain name.

Default

N/A.

Usage Guidelines

The default domain name will be used to create a fully qualified host name when a domain name is not specified. For example, if the default domain name is set to "food.com" then when a command like "ping dog" is entered, the ping will actually be executed as "ping dog.food.com".

Example

The following command configures the default domain name for the server:

```
configure dns-client default-domain xyz_inc.com
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure dns-client delete

```
configure dns-client delete [domain-suffix <domain_name> | name-server
<ip_address> {vr <vr_name>}]
```

Description

Deletes a domain suffix from the domain suffix list or a name server from the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies deleting a domain suffix.
domain_name	Specifies a domain name.
name-server	Specifies deleting a name server.
ip_address	Specifies an IP address for the name server.
vr	Specifies deleting a virtual router.
	NOTE: User-created VRs are not available on the BlackDiamond 8800 series switch, SummitStack, or the Summit family of switches.
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Specifying a domain suffix removes an entry from the domain suffix list. If the deleted item was not the last entry in the list, all items that had been added later are moved up in the list. If no entries in the list match the domain name specified, an error message will be displayed.

The default value for the virtual router used by the DNS client option is VR-Default.

Examples

The following command deletes a domain name from the domain suffix list:

```
configure dns-client delete domain-suffix xyz_inc.com
```

The following command removes a DNS server from the list:

```
configure dns-client delete name-server 10.1.2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure failsafe-account

```
configure failsafe-account
```

Description

Configures a name and password for the failsafe account.

Syntax Description

This command has no arguments or variables.

Default

The failsafe account is always configured.

Usage Guidelines

The failsafe account is the account of last resort to access your switch.

You will be prompted for the failsafe account name, and prompted twice to specify the password for the account. For example:

```
BD-10808.1 # configure failsafe-account
enter failsafe user name: blue5green
enter failsafe password:
enter password again:
BD-10808.2
```

The failsafe account is immediately saved to NVRAM on all MSMs or active nodes in a stack.



NOTE

The information that you use to configure the failsafe account cannot be recovered by Extreme Networks. Technical support cannot retrieve passwords or account names for this account. Protect this information carefully.

Once you enter the failsafe account name, you are prompted to enter the password. Once you successfully log in to the failsafe account, you are logged in to an admin-level account.

Example

The following command changes the failsafe account:

```
configure failsafe-account
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure idletimeout

```
configure idletimeout <minutes>
```

Description

Configures the time-out for idle console, SSH2, and Telnet sessions.

Syntax Description

minutes	Specifies the time-out interval, in minutes. Range is 1 to 240 (1 minute to 4 hours).

Default

The default time-out is 20 minutes.

Usage Guidelines

This command configures the length of time the switch will wait before disconnecting idle console, SSH2, or Telnet sessions. The idletimeout feature must be enabled for this command to have an effect (the idletimeout feature is enabled by default).

Example

The following command sets the time-out for idle login and console sessions to 10 minutes:

```
configure idletimeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure safe-default-script

```
configure safe-default-script
```

Description

Allows you to change management access to your device and to enhance security. Runs an interactive script that prompts you to choose to enable or disable SNMP, Telnet, and enabled ports. Refer to [“Safe Defaults Setup Method” on page 53](#) in the *Extreme Networks XOS Concepts Guide* for complete information on the safe default mode.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Once you issue this command, the system presents you with the following interactive script:

Telnet is enabled by default. Telnet is unencrypted and has been the target of security exploits in the past.

Would you like to disable Telnet? [y/N]:

SNMP access is enabled by default. SNMP uses no encryption, SNMPv3 can be configured to eliminate this problem.

Would you like to disable SNMP? [y/N]:

All ports are enabled by default. In some secure applications, it maybe more desirable for the ports to be turned off.

Would you like unconfigured ports to be turned off by default? [y/N]:

Changing the default failsafe account username and password is highly recommended. If you choose to do so, please remember the username and password as this information cannot be recovered by Extreme Networks. Would you like to change the failsafe account username and password now? [y/N]:

Would you like to permit failsafe account access via the management port? [y/N]:

Since you have chosen less secure management methods, please remember to increase the security of your network by taking the following actions:

- * change your admin password
- * change your failsafe account username and password

- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

Example

The following command reruns the interactive script to configure management access:

```
configure safe-default-script
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure time

```
configure time <month> <day> <year> <hour> <min> <sec>
```

Description

Configures the system date and time.

Syntax Description

month	Specifies the month. The range is 1-12.
day	Specifies the day of the month. The range is 1-31.
year	Specifies the year in the YYYY format. The range is 2003 to 2036.
hour	Specifies the hour of the day. The range is 0 (midnight) to 23 (11 pm).
min	Specifies the minute. The range is 0-59.
sec	Specifies the second. The range is 0-59.

Default

N/A.

Usage Guidelines

The format for the system date and time is as follows:

```
mm dd yyyy hh mm ss
```

The time uses a 24-hour clock format. You cannot set the year earlier than 2003 or past 2036. You have the choice of inputting the entire time/date string. If you provide one item at a time and press [Tab], the screen prompts you for the next item. Press <cr> to complete the input.

Example

The following command configures a system date of February 15, 2002 and a system time of 8:42 AM and 55 seconds:

```
configure time 02 15 2002 08 42 55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure timezone

```
configure timezone {name <tz_name>} <GMT_offset>
{autodst {name <dst_timezone_ID>} {<dst_offset>}}
{begins [every <floatingday> | on <absoluteday>] {at <time_of_day>}}
{ends [every <floatingday> | on <absoluteday>] {at <time_of_day>}}}
| noautodst}
```

Description

Configures the Greenwich Mean Time (GMT) offset and Daylight Saving Time (DST) preference.

Syntax Description

tz_name	Specifies an optional name for this timezone specification. May be up to six characters in length. The default is an empty string.
GMT_offset	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
autodst	Enables automatic Daylight Saving Time.
dst-timezone-ID	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
dst_offset	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.
floatingday	Specifies the day, week, and month of the year to begin or end DST each year. Format is: <week> <day> <month> where: • <week> is specified as [first second third fourth last] or 1-5. • <day> is specified as [sunday monday tuesday wednesday thursday friday saturday] or 1-7 (where 1 is Sunday). • <month> is specified as [january february march april may june july august september october november december] or 1-12. Default for beginning is first sunday april; default for ending is last sunday october.
absoluteday	Specifies a specific day of a specific year on which to begin or end DST. Format is: <month> <day> <year> where: • <month> is specified as 1-12. • <day> is specified as 1-31. • <year> is specified as 2003-2035. The year must be the same for the begin and end dates.
time_of_day	Specifies the time of day to begin or end Daylight Saving Time. May be specified as an hour (0-23) or as hour:minutes. Default is 2:00.
noautodst	Disables automatic Daylight Saving Time.

Default

Autodst, beginning every first Sunday in April, and ending every last Sunday in October.

Usage Guidelines

Network Time Protocol (NTP) server updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographic location.

The `gmt_offset` is specified in +/- minutes from the GMT time.

Automatic DST changes can be enabled or disabled. The default configuration, where DST begins on the first Sunday in April at 2:00 AM and ends the last Sunday in October at 2:00 AM, applies to most of North America, and can be configured with the following syntax:

```
configure timezone <gmt_offset> autodst.
```

The starting and ending date and time for DST may be specified, as these vary in time zones around the world.

- Use the `every` keyword to specify a year-after-year repeating set of dates (for example, the last Sunday in March every year)
- Use the `on` keyword to specify a non-repeating, specific date for the specified year. If you use this option, you will need to specify the command again every year.
- The `begins` specification defaults to `every first sunday april`.
- The `ends` specification defaults to `every last sunday october`.
- The `ends` date may occur earlier in the year than the `begins` date. This will be the case for countries in the Southern Hemisphere.
- If you specify only the starting or ending time (not both) the one you leave unspecified will be reset to its default.
- The `time_of_day` specification defaults to `2:00`.
- The timezone IDs are optional. They are used only in the display of timezone configuration information in the `show switch` command.

To disable automatic DST changes, re-specify the GMT offset using the `noautodst` option:

```
configure timezone <gmt_offset> noautodst.
```

NTP updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location. [Table 6](#) describes the GMT offsets.

Table 6: Greenwich Mean Time offsets

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+0:00	+0	GMT - Greenwich Mean UT or UTC - Universal (Coordinated) WET - Western European	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
-1:00	-60	WAT - West Africa	Cape Verde Islands
-2:00	-120	AT - Azores	Azores
-3:00	-180		Brasilia, Brazil; Buenos Aires, Argentina; Georgetown, Guyana;
-4:00	-240	AST - Atlantic Standard	Caracas; La Paz
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, Trevor City, MI USA
-6:00	-360	CST - Central Standard	Mexico City, Mexico

Table 6: Greenwich Mean Time offsets (Continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
-7:00	-420	MST - Mountain Standard	Saskatchewan, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA, Cupertino, CA, Seattle, WA USA
-9:00	-540	YST - Yukon Standard	
-10:00	-600	AHST - Alaska-Hawaii Standard	
		CAT - Central Alaska	
		HST - Hawaii Standard	
-11:00	-660	NT - Nome	
-12:00	-720	IDLW - International Date Line West	
+1:00	+60	CET - Central European	Paris, France; Berlin, Germany;
		FWT - French Winter	Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain;
		MET - Middle European	Rome, Italy; Bern, Switzerland;
		MEWT - Middle European Winter	Stockholm, Sweden; Oslo, Norway
		SWT - Swedish Winter	
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	
+5:30	+330	IST - India Standard Time	New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6 - Russia Zone 5	
+7:00	+420	WAST - West Australian Standard	
+8:00	+480	CCT - China Coast, Russia Zone 7	
+9:00	+540	JST - Japan Standard, Russia Zone 8	
+10:00	+600	EAST - East Australian Standard	
		GST - Guam Standard	
		Russia Zone 9	
+11:00	+660		
+12:00	+720	IDLE - International Date Line East	Wellington, New Zealand; Fiji, Marshall Islands
		NZST - New Zealand Standard	
		NZT - New Zealand	

Example

The following command configures GMT offset for Mexico City, Mexico and disables automatic DST:

```
configure timezone -360 noautodst
```

The following four commands are equivalent, and configure the GMT offset and automatic DST adjustment for the US Eastern timezone, with an optional timezone ID of EST:

```
configure timezone name EST -300 autodst name EDT 60 begins every first sunday april  
at 2:00 ends every last sunday october at 2:00
```

```
configure timezone name EST -300 autodst name EDT 60 begins every 1 1 4 at 2:00 ends  
every 5 1 10 at 2:00
```

```
configure timezone name EST -300 autodst name EDT
```

```
configure timezone -300 autodst
```

The following command configures the GMT offset and automatic DST adjustment for the Middle European timezone, with the optional timezone ID of MET:

```
configure timezone name MET 60 autodst name MDT begins every last sunday march at 1  
ends every last sunday october at 1
```

The following command configures the GMT offset and automatic DST adjustment for New Zealand. The ending date must be configured each year because it occurs on the first Sunday on or after March 5:

```
configure timezone name NZST 720 autodst name NZDT 60 begins every first sunday  
october at 2 ends on 3/16/2002 at 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create account

```
create account [admin | user] <account-name> {encrypted <password>}
```

Description

Creates a new user account.

Syntax Description

admin	Specifies an access level for account type admin.
user	Specifies an access level for account type user.
account-name	Specifies a new user account name. See “Usage Guidelines” for more information.
encrypted	Specifies the encrypted option.
password	Specifies a user password. See “Usage Guidelines” for more information.

Default

By default, the switch is configured with two accounts with the access levels shown in [Table 7](#).

Table 7: User account levels

Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	This user can view (but not change) all manageable parameters, with the following exceptions: • This user cannot view the user account database. • This user cannot view the SNMP community strings. • This user cannot view SSL settings. This user has access to the ping command.

You can use the default names (*admin* and *user*), or you can create new names and passwords for the accounts. Default accounts do not have passwords assigned to them.

Usage Guidelines

The switch can have a total of 16 user accounts. The system must have one administrator account.

When you use the `encrypted` keyword, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should not use the encrypted option and should enter the password in plain text. The encrypted option is used by the switch after encrypting the plain text password. The encrypted option should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.

The system prompts you to specify a password after you enter this command and to reenter the password. If you do not want a password associated with the specified account, press *Enter twice*.

You must have administrator privileges to change passwords for accounts other than your own. User names and passwords are case-sensitive. User account names must have a minimum of 1 character and can have a maximum of 32 characters. Passwords must have a minimum of 0 characters and can have a maximum of 32 characters. Passwords are case-sensitive; user names are not case-sensitive.

**NOTE**

If the account is configured to require a specific password format, the minimum is 8 characters. See [configure account password-policy char-validation](#) for more information.

Example

The following command creates a new account named John2 with administrator privileges:

```
create account admin John2
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted option was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

delete account

```
delete account <name>
```

Description

Deletes a specified user account.

Syntax Description

name	Specifies a user account name.

Default

N/A.

Usage Guidelines

Use the `show accounts` command to determine which account you want to delete from the system. The `show accounts` output displays the following information in a tabular format:

- The user name
- Access information associated with each user
- User login information
- Session information

Depending on the software version running on your switch and the type of switch you have, additional account information may be displayed.

You must have administrator privileges to delete a user account. The system must have one administrator account; the command will fail if an attempt is made to delete the last administrator account on the system.

To ensure security, change the password on the default account, but do not delete it. The changed password will remain intact through configuration uploads and downloads.

If you must delete the default account, first create another administrator-level account.

Example

The following command deletes account John2:

```
delete account John2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable cli space-completion

```
disable cli space-completion
```

Description

Disables the ExtremeXOS feature that completes a command automatically with the spacebar. If you disable this feature, the [Tab] key can still be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables using the spacebar to automatically complete a command:

```
disable cli space-completion
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable clipaging

```
disable clipaging
```

Description

Disables pausing at the end of each show screen.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most `show` command output will pause when the display reaches the end of a page. This command disables the pause mechanism and allows the display to print continuously to the screen.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

Example

The following command disables clipaging and allows you to print continuously to the screen:

```
disable clipaging
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable idletimeout

```
disable idletimeout
```

Description

Disables the timer that disconnects idle sessions from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled. Timeout 20 minutes.

Usage Guidelines

When idle time-outs are disabled, console sessions remain open until the switch is rebooted or until you logoff. Telnet sessions remain open until you close the Telnet client.

If you have an SSH2 session and disable the idle timer, the SSH2 connection times out after 61 minutes of inactivity.

To view the status of idle time-outs on the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enable/disable state for idle time-outs.

Example

The following command disables the timer that disconnects all sessions to the switch:

```
disable idletimeout
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable cli space-completion

```
enable cli space-completion
```

Description

Enables the ExtremeXOS feature that completes a command automatically with the spacebar. The [Tab] key can also be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables using the spacebar to automatically complete a command:

```
enable cli space-completion
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable clipaging

```
enable clipaging
```

Description

Enables the pause mechanism and does not allow the display to print continuously to the screen.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most `show` command output will pause when the display reaches the end of a page.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

If CLI paging is enabled and you use the `show tech` command to diagnose system technical problems, the CLI paging feature is disabled.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

Example

The following command enables clipaging and does not allow the display to print continuously to the screen:

```
enable clipaging
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable idletimeout

```
enable idletimeout
```

Description

Enables a timer that disconnects Telnet, SSH2, and console sessions after a period of inactivity (20 minutes is default).

Syntax Description

This command has no arguments or variables.

Default

Enabled. Timeout 20 minutes.

Usage Guidelines

You can use this command to ensure that a Telnet, Secure Shell (SSH2), or console session is disconnected if it has been idle for the required length of time. This ensures that there are no hanging connections.

To change the period of inactivity that triggers the timeout for a Telnet, SSH2, or console session, use the `configure timezone` command.

To view the status of idle timeouts on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle timeouts. You can configure the length of the timeout interval.

Example

The following command enables a timer that disconnects any Telnet, SSH2, and console sessions after 20 minutes of inactivity:

```
enable idletimeout
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable license software

```
enable license {software} <key>
```

Description

Enables software license or feature pack that allows you to use advanced features.

Syntax Description

key	Specifies your hexadecimal license key in format xxxx-xxxx-xxxx-xxxx-xxxx.

Default

N/A

Usage Guidelines

The software license levels that apply to ExtremeXOS version 11.6 are shown in [Table 8](#).

Table 8: Licensing requirements for ExtremeXOS version 11.6

License level	BlackDiamond 10808 switch	BlackDiamond 12800 series switch	Summit X450e switches	Summit X450 and X450a switches	BlackDiamond 8800 series switch
Edge	Ships (as part of Core license)		Ships	Ships	Ships
Advanced Edge	Ships (as part of Core license)		Requires license	Ships	Ships
Core	Ships	- MSM-5 ships - MSM-5R ships	N/A	Requires license	Requires license
Advanced Core	- MSM 1—N/A - MSM1XL—Ships	N/A	N/A	N/A	N/A

With ExtremeXOS software version 11.2, BGP functionality moved from the Advanced Core license level to the Core license. You have BGP functionality with a Core license.

To obtain a software license, specify the `key` in the format `xxxx-xxxx-xxxx-xxxx-xxxx`.

You obtain the software license key (or feature pack key) either by ordering it from the factory or by obtaining a license voucher from your Extreme Networks supplier. You can obtain a regular software license or a trial software license, which allows you use of the license for either 30, 60 or 90 days; you cannot downgrade software licenses.

The voucher contains all the necessary information on the software license, whether regular or trial, and number of days for trial software license.

After you enable the software license or feature pack by entering the software key, the system returns a message that you either successfully or unsuccessfully set the license.

Once you enable the software license (or if you do not use the correct key, attempt to downgrade the license, or already installed the software license) you see one of the following messages:

```
Enabled license successfully.
Error: Unable to set license using supplied key.
Error: Unable to set license - downgrade of licenses is not supported.
Error: Unable to set license - license is already enabled.
Error: Unable to set license - trial license already enabled.
```

If you enable a trial license, the system generates a daily message showing the number of days until expiry.

Once installed (or enabled), the software license goes with the switch chassis itself (not with the MSM module on modular switches). A software license must be installed separately on each Summit X450 or X250 series switch, whether or not it is operating in a SummitStack. The software license information is stored in EEPROM in the modular switches and NVRAM on the stand-alone switches and SummitStack; the information persists through reboots, software upgrades, power outages, and reconfigurations.

If you attempt to execute a command and you do not either have the required software license or have reached the limits defined by the current software license level, the system returns one of the following messages:

```
Error: This command cannot be executed at the current license level.
Error: You have reached the maximum limit for this feature at this license level.
```

If you attempt to execute a command and you do not have the required feature pack, the system also returns a message.

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

To view the type of software license you are currently running on the switch, use the `show licenses` command. The license key number is not displayed, but the type of software license is displayed in the `show licenses` output. This command can be run on any node in a SummitStack, regardless of its node role (Master, Standby, or Backup).

Example

The following command enables a software license on the switch:

```
enable license 2d5e-0e84-e87d-c3fe-bfff
```

History

This command was first available in ExtremeXOS 11.1.

The `software` parameter was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable license file

```
enable license file <filename>
```

Description

Enables the text file that applies software licenses and feature packs licenses to more than one switch at a time.

Syntax Description

filename	Specifies the filename that you download onto the switch using TFTP; the file extension is .xlic.

Default

N/A

Usage Guidelines

You download the license file to the switch using TFTP or SCP. The file name extension for this file is <xlic>; for example, you may see a file named systemlic.xlic.

Using this file, you enable the software and feature pack licenses for more than one switch simultaneously. The file can contain licenses for some or all of the Extreme Networks switches that the customer owns. During upload, only those license keys destined for the specific switch are used to attempt enabling the licenses. The license file is a text file that has the switch serial number, software license type, and license key; it is removed from the switch after the licenses are enabled.

After you enable the license file, the system returns one or more of the following messages:

```
Enabled license successfully.
Error: Unable to set license <license_name> using supplied key.
Error: Unable to set license <license_name> - downgrade of licenses is not supported.
Error: Unable to set license <license_name> - license is already enabled.
Error: Unable to set license <license_name> - trial license already enabled.
```

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

For Summit X450 and X250 switches, the command must be executed on each node, whether or not it is operating in a SummitStack.

Example

The following command enables a license file on the specified Extreme Networks switches:

```
enable license file santaclara.xlic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

history

history

Description

Displays a list of all the commands entered on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

ExtremeXOS “remembers” all the commands you entered on the switch. Use the `history` command to display a list of these commands.

Example

The following command displays all the commands entered on the switch:

```
history
```

If you use a command more than once, consecutively, the history will only list the first instance.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

ping

```
ping {count <count> {start-size <start-size>} | continuous {start-size
<start-size>} | {start-size <start-size> {end-size <end-size>}}} {udp}
{dont-fragment} {ttl <ttl>} {tos <tos>} {interval <interval>} {vr <vrid>}
{ipv4 <host> | ipv6 <host>} {from} {with record-route}
```

Description

Enables you to send User Datagram Protocol (UDP) or Internet Control Message Protocol (ICMP) echo messages or to a remote IP device.

Syntax Description

count	Specifies the number of ping requests to send.
start-size	Specifies the size, in bytes, of the packet to be sent, or the starting size if incremental packets are to be sent.
continuous	Specifies that UDP or ICMP echo messages to be sent continuously. This option can be interrupted by pressing [Ctrl] + C.
end-size	Specifies an end size for packets to be sent.
udp	Specifies that the ping request should use UDP instead of ICMP.
dont-fragment	Sets the IP to not fragment the bit.
ttl	Sets the TTL value.
tos	Sets the TOS value.
interval	Sets the time interval between sending out ping requests.
vr	Specifies the virtual route to use for sending out the echo message. If not specified, VR-Default is used. NOTE: User-created VRs are not available on the BlackDiamond 8800 series switch, SummitStack, or the Summit family of switches.
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport. NOTE: If you are contacting an IPv6 link local address, you must specify the VLAN you are sending the message from: ping <ipv6> <link-local address> %<vlan_name> <host>.
host	Specifies a host name or IP address (either v4 or v6).
from	Uses the specified source address. If not specified, the address of the transmitting interface is used.
with record-route	Sets the traceroute information.

Default

N/A.

Usage Guidelines

**NOTE**

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

The `ping` command is used to test for connectivity to a specific host.

You use the `ipv6` variable to ping an IPv6 host by generating an ICMPv6 echo request message and sending the message to the specified address. If you are contacting an IPv6 link local address, you must specify the VLAN you sending the message from, as shown in the following example (you must include the % sign): `ping <ipv6> <link-local address> %<vlan_name> <host>`.

**NOTE**

You must specify the target's IPv6 address to use this functionality. DNS for IPv6 is not supported, so the `ping` resolves to an IPv4 address if you specify a host name.

The `ping` command is available for both the user and administrator privilege level.

Example

The following command enables continuous ICMP echo messages to be sent to a remote host:

```
ping continuous 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 variable was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

reboot

For modular switches:

```
reboot {time <month> <day> <year> <hour> <min> <sec>} {cancel} {msm
<slot_id>} {slot <slot-number> | node-address <node-address> | stack-
topology {as-standby} }
```

For Summit X450 and X250 series switches and SummitStack:

```
reboot {[time <mon> <day> <year> <hour> <min> <sec>] | cancel} {slot <slot-
number> | node-address <node-address> | stack-topology {as-standby}}
```

Description

Reboots the switch, SummitStack, or the module in the specified slot at a specified date and time.

Syntax Description

time	Specifies a reboot date in mm dd yyyy format and reboot time in hh mm ss format.
cancel	Cancels a previously scheduled reboot.
msm	Specifies rebooting the MSM module. NOTE: This variable is available only on modular switches.
slot_id	Specifies the slot--A or B--for an MSM module. NOTE: This variable is available only on modular switches.
slot-number	Specifies the slot number currently being used by the active stack node that is to be rebooted (SummitStack only).
node-address	Specifies the MAC address of the SummitStack node to be rebooted (SummitStack and Summit Family of switches only).
stack-topology	Specifies that the entire SummitStack is to be rebooted whether or not nodes are active (SummitStack and Summit Family of switches only).
as-standby	Specifies that all stack nodes that are to be rebooted are to operate as if configured to not be master-capable (SummitStack and Summit Family of switches only).

Default

N/A.

Usage Guidelines

If you do not specify a reboot time, the switch will reboot immediately following the command, and any previously scheduled reboots are cancelled. Prior to rebooting, the switch returns the following message:

```
Do you want to save configuration changes to primary and reboot?
(y - save and reboot, n - reboot without save, <cr> - cancel command)
```

To cancel a previously scheduled reboot, use the `cancel` option.

Modular switches only. The modules that can be rebooted are switch fabric MSM modules.

BlackDiamond 8800 series switch only. On the BlackDiamond 8800 series switch, if your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a sharp object into the “A” and “R” holes on the MSM and applying slight pressure. Refer to *Extreme Networks Consolidated XOS Hardware Installation Guide* for information on the MSM.

The `reboot MSM` option on the BlackDiamond 8800 series switch affects the entire module.

SummitStack and Summit Family of switches only. The reboot command used without any parameters on the master node reboots all members of the same active topology to which the master node belongs. This version can only be used on the master node.

The `reboot slot <slot-number>` command can be used on any active node. The command will reboot the active node that is currently using the specified slot number in the same active topology as the issuing node. This variation cannot be used on a node that is not in stacking mode.

The `reboot node-address <node-address>` command can be used on any node whether or not the node is in stacking mode. It will reboot the node whose MAC address is supplied.

The `reboot stack-topology {as-standby}` command reboots every node in the stack topology. The command can be issued from any node whether or not the node is in stacking mode. If the `as-standby` option is used, every node in the stack topology restarts with master-capability disabled. This option is useful when manually resolving a dual master situation.

Example

The following command reboots the switch at 8:00 AM on April 15, 2005:

```
reboot time 04 15 2005 08 00 00
```

History

This command was first available in ExtremeXOS 10.1.

The alternate BootROM image was added in ExtremeXOS 11.1.

The slot, node-address, stack-topology, and as-standby options were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

The alternate BootROM is available only on the BlackDiamond 8800 series switch.

The node-address, stack-topology and as-standby are available only in SummitStack and the Summit family of switches.

The slot is available only on switches operating in a SummitStack.

show accounts

```
show accounts
```

Description

Displays user account information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You need to create a user account using the `create account` command before you can display user account information.

To view the accounts that have been created, you must have administrator privileges.

The `show accounts` command displays the following information in a tabular format:

- User Name—The name of the user. This list displays all of the users who have access to the switch.
- Access—This may be listed as R/W for read/write or RO for read only.
- Login OK—The number of logins that are okay.
- Failed—The number of failed logins.
- Accounts locked out—Account configured to be locked out after 3 consecutive failed login attempts (using the `configure account password-policy lockout-on-login-failures` command).



NOTE

This command does not show the failsafe account.

Example

The following command displays user account information on the switch:

```
show accounts pppuser
```

Output from this command looks similar to the following:

User Name	Access	LoginOK	Failed
admin	R/W	3	1
user	RO	0	0
dbackman	R/W	0	0
ron*	RO	0	0

```
nocteam    RO          0          0
-----
(*) - Account locked
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show accounts password-policy

```
show accounts password-policy
```

Description

Displays password policy information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

To view the password management information, you must have administrator privileges.

The `show accounts password-policy` command displays the following information in a tabular format:

- Global password management parameters applied to new accounts upon creation:
 - Maximum age—The maximum number of days for the passwords to remain valid.
 - History limit—The number of previous password that the switch scans prior to validating a new password.
 - Minimum length—The minimum number of characters in passwords.
 - Character validation—The passwords must be in the specific format required by the `configure account password-policy char-validation` command.
 - Lockout on login failures—If enabled, the system locks out users after 3 failed login attempts.
 - Accounts locked out—Number of accounts locked out.
- User Name—The name of the user. This list displays all of the users who have access to the switch.
- Password Expiry Date—Date the password for this account expires; may be blank.
- Password Max. age—The number of days originally allowed to passwords on this account; may show None.
- Password Min. length—The minimum number of characters required for passwords on this account; may show None.
- Password History Limit—The number of previous passwords the system scans to disallow duplication on this account; may show None.

Example

The following command displays the password management parameters configured for each account on the switch:

```
show accounts password-policy
```

Output from this command looks similar to the following:

```

-----
Accounts global configuration(applied to new accounts on creation)
-----
Password Max. age           : None
Password History limit      : None
Password Min. length        : None
Password Character Validation : Disabled
Accts. lockout on login failures: Disabled
Accounts locked out         : No
-----

```

User Name	Password Expiry Date	Password Max. age	Password Min. len	Password History Limit	Flags
admin		None	None	None	---
user		None	None	None	---
test	Apr-17-2005	12	32	9	C--

```

-----
Flags: (C) Password character validation enabled, (L) Account locked out
       (l) Account lockout on login failures enabled
-----

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show banner

```
show banner
```

Description

Displays the user-configured banner string.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed before the login prompt.

Example

The following command displays the switch banner:

```
show banner
```

Output from this command varies depending on your configuration; the following is one example:

```
Extreme Networks Summit X450 Switch
#####
  Unauthorized Access is strictly prohibited.
  Violators will be prosecuted
#####
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show dns-client

```
show dns-client
```

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the DNS configuration:

```
show dns-client
```

Output from this command looks similar to the following:

```
Number of domain suffixes: 2
Domain Suffix 1:          njudah.local
Domain Suffix 2:          dbackman.com
Number of name servers: 2
Name Server 1:  172.17.1.104
Name Server 2:  172.17.1.123
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show licenses

```
show licenses
```

Description

Displays current software license level and feature packs enabled on your switches.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The command displays information on the software license level and feature packs enabled on the switch, including the trial license and days left to expiry.

To see the license information on a SummitStack node, run the command while logged into that node.



NOTE

Refer to the specific chapter that discusses each feature of the ExtremeXOS Concepts Guide to determine if a license is required for some functionality. If not noted, all functionality is available, and license is not required.

Example

The following command displays the license level configuration:

```
show licenses
```

Output from this command looks similar to the following:

```
BD-10808.4 # show licenses
Enabled License Level:
    Core
Enabled Feature Packs:
    MPLS
```

On a SummitStack, the output from this command looks similar to the following:

```
Slot-3 Stack.12 > show licenses
Enabled License Level:
    Advanced Edge
Enabled Feature Packs:
    None
Effective License Level:
    Edge
```

History

This command was first available in ExtremeXOS 11.1.

The information on enabled feature packs was added in ExtremeXOS 11.4.

The information on the trial licenses was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show switch

```
show switch {detail}
```

Description

Displays the current switch information.

On a SummitStack, this command displays the Master and Backup node information if executed on the Master, and displays the current node and the Master node information if executed on any other node.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show switch` command displays:

- sysName, sysLocation, sysContact
- MAC address
- System type
- System health check
- Recovery mode
- Watchdog state
- Current date, time, system boot time, and time zone configuration
- Any scheduled reboot information
- System up time
- Master and Backup information (available only on modular switches and SummitStack)
- Current state (available only on stand-alone switches)
 - OPERATIONAL
 - OPERATIONAL (OverHeat)
 - FAILED
- Software image information (primary/secondary image and version)
- Configuration information (primary/secondary configuration and version)

This information may be useful for your technical support representative if you have a problem.

On a SummitStack, the System UpTime may be useful when manually resolving the dual master situation. For more information, see the “Eliminating a Dual Master Situation Manually” section in the *ExtremeXOS Concepts Guide*.

Depending on the software version running on your switch, additional or different switch information may be displayed.

On a stack the following additional information will be available:

- System Type
- System UpTime
- Details of Master and Backup, or current node and Master

Example

The following command displays current switch information:

```
show switch
```

Output from this command on the modular switches looks similar to the following:

```
SysName:          Aspen15
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1D:00:C0
System Type:      BD-8810

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled

Current Time:     Fri Feb 13 02:25:24 1925
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Wed Feb 11 21:39:56 1925
Boot Count:       159
Next Reboot:      None scheduled
System UpTime:    1 day 4 hours 45 minutes 28 seconds

Slot:             MSM-A *                MSM-B
                  -----                -----
Current State:    MASTER                  BACKUP (In Sync)

Image Selected:   secondary                secondary
Image Booted:     primary                  primary
Primary ver:      12.0.0.4                  12.0.0.4
Secondary ver:    12.0.0.4                  12.0.0.4

Config Selected:  primary.cfg              primary.cfg
Config Booted:    primary.cfg              primary.cfg

primary.cfg       Created by ExtremeXOS version 11.6.0.30
                  574246 bytes saved on Wed Jul 30 19:39:55 1924
```

Output from this command on the stand-alone Summit family switch looks similar to the following:

```
SysName:          X450a-24tdc
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:26:6B:EC
```

```

System Type:      X450a-24tdc

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled

Current Time:     Wed Apr 25 21:17:18 2012
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Wed Apr 25 21:13:54 2012
Boot Count:       951
Next Reboot:      None scheduled
System UpTime:    3 minutes 24 seconds

Current State:    OPERATIONAL
Image Selected:   secondary
Image Booted:     secondary
Primary ver:      12.0.0.4
Secondary ver:    12.0.0.4

Config Selected:  primary.cfg
Config Booted:    Factory Default

primary.cfg       Created by ExtremeXOS version 12.0.0.4
                  156281 bytes saved on Mon Apr 23 17:10:11 2012

```

The `show switch detail` command displays the same information shown above.

Output from this command on a stack looks similar to the following:

```

SysName:          Stack
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       02:04:96:27:B7:41
System Type:      X450e-24p (Stack)

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled

Current Time:     Tue Jan 30 14:22:41 2007
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Mon Jan 29 21:51:38 2007
Boot Count:       317
Next Reboot:      None scheduled
System UpTime:    16 hours 31 minutes 3 seconds

Slot:             Slot-4 *                               Slot-5
                  -----                               -----
Current State:    MASTER                                  BACKUP (In Sync)

Image Selected:   secondary                               secondary
Image Booted:     secondary                               secondary
Primary ver:      12.0.0.10                               12.0.0.10
Secondary ver:    12.0.0.13                               12.0.0.13

Config Selected:  primary.cfg

```

```
Config Booted:    primary.cfg

primary.cfg       Created by ExtremeXOS version 12.0.0.10
                  139108 bytes saved on Fri Jan 26 22:56:40 2007
```

History

This command was first available in ExtremeXOS 10.1.

This command was updated to support stacking in ExtremeXOS 12.0 and the System Type was added to the output from this version.

Platform Availability

This command is available on all platforms.

traceroute

```
traceroute {vr <vrid>} {ipv4 <host>} {ipv6 <host>} {ttl <number>} {from
<from>} {[port <port>] | icmp}
```

Description

Enables you to trace the routed path between the switch and a destination endstation.

Syntax Description

vr	Specifies a virtual router. NOTE: User-created VRs are not available on the BlackDiamond 8800 series switch, SummitStack, or the Summit family of switches.
vrid	Specifies which virtual router. NOTE: User-created VRs are not available on the BlackDiamond 8800 series switch, SummitStack, or the Summit family of switches.
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport.
host	Specifies the host of the destination endstation.
ttl <number>	Configures the switch to trace up to the time-to-live number of the switch.
from <from>	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
port <port>	Specifies the UDP port number.
icmp	Configures the switch to send ICMP echo messages to trace the routed path between the switch and a destination endstation.

Default

N/A.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created virtual routers.

Each router along the path is displayed.

Beginning with ExtremeXOS 11.2, you can trace the routed path between the switch and the given IPv6 address.



NOTE

You must specify the target's IPv6 address to use this functionality. DNS for IPv6 is not supported, so the traceroute resolves to an IPv4 address if you specify a host name.

Example

The following command enables the traceroute function to a destination of 123.45.67.8:

```
traceroute 123.45.67.8
```

The following is sample output that displays when the traceroute fails:

```
traceroute to 10.209.10.37, 30 hops max
 1  0.0.0.0                                * !u          * !u          * !u

--- Packet Response/Error Flags ---
(*) No response, (!N) ICMP network unreachable, (!H) ICMP host unreachable,
(!P) ICMP protocol unreachable, (!F) ICMP fragmentation needed,
(!S) ICMP source route failed, (!u) Transmit error, network unreachable,
(!f) Transmit error, fragmentation needed, (!t) General transmit error
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 variable was added in ExtremeXOS 11.2.

The display when the command fails was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

3

Commands for Managing the Switch

This chapter describes commands for:

- Configuring Simple Network Management Protocol (SNMP) parameters on the switch
- Managing the switch using Telnet
- Transferring files using the Trivial File Transfer Protocol (TFTP)
- Configuring system redundancy
- Displaying power management statistics on the switch
- Configuring Simple Network Time Protocol (SNTP) parameters on the switch

SNMP

Any network manager running the Simple Network Management Protocol (SNMP) can manage the switch, if the Management Information Base (MIB) is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

The following SNMP parameters can be configured on the switch:

- Authorized trap receivers—An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers. Entries in this list can be created, modified, and deleted using the RMON2 trapDestTable MIB table, as described in RFC 2021, and the SNMPv3 tables.
- Authorized managers—An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask.
- Community strings—The community strings allow a simple method of authentication between the switch and the remote network manager. The default read-only community string is *public*. The default read-write community string is *private*. The community strings for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps.
- System contact (optional)—The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch.
- System name (optional)—The system name enables you to enter a name that you have assigned to this switch. The default name is the model name of the switch (for example, BD-1.2).
- System location (optional)—Using the system location field, you can find the location of the switch.



NOTE

If you specify volatile storage when configuring SNMP parameters, that configuration is not saved across a switch reboot.

Telnet

Telnet allows you to access the switch remotely using TCP/IP through one of the switch ports or a workstation with a Telnet facility. If you access the switch via Telnet, you will use the command line interface (CLI) to manage the switch and modify switch configurations.

TFTP

ExtremeXOS supports the Trivial File Transfer Protocol (TFTP) based on RFC 1350. TFTP is a method used to transfer files from one network device to another. The ExtremeXOS TFTP client is a command line application used to contact an external TFTP server on the network. For example, ExtremeXOS uses TFTP to download software image files, switch configuration files, and access control lists (ACLs) from a server on the network to the switch.

System Redundancy with Dual MSMs Installed—Modular Switches Only

If you install two MSMs in a modular switch, one assumes the role of primary and the other assumes the role of backup. The primary MSM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary also keeps synchronized with the backup MSM in case the backup MSM needs to take over the management functions if the primary MSM fails.

Power Supply Management

On modular switches, ExtremeXOS monitors and manages power consumption on the switch by periodically checking the power supply units (PSUs) and testing them for failures. To determine the health of the PSU, ExtremeXOS checks the voltage, current, and temperature of the PSU.

The power management capability of ExtremeXOS:

- Protects the system from overload conditions
- Monitors all installed PSUs, even installed PSUs that are disabled
- Enables and disables PSUs as required
- Powers up or down I/O modules based on available power and required power resources
- Logs power resource changes, including power budget, total available power, redundancy, and so on
- Detects and isolates faulty PSUs

On the Summit family switches, ExtremeXOS reports when the PSU has power or has failed. The Summit family switches support an internal power supply with a range of 90V to 240V AC power as well as an external redundant power supply. The Extreme Networks External Power System (EPS) allows you to add a redundant power supply to the Summit family switch to protect against a power supply failure. The EPS consists of a tray (EPS-T) that holds one or two EPS-160 power supplies. The

EPS-160 provides 100V to 240V AC power. Each EPS-160 power supply provides one-to-one redundancy to an attached Summit family switch.

Simple Network Time Protocol

ExtremeXOS supports the client portion of the Simple Network Time Protocol (SNTP) Version 3 based on RFC1769. SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. When enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean time (GMT) offset and the use of Daylight Saving Time.

configure node priority

```
configure node slot <slot_id> priority <node_pri>
```

Description

Configures the priority of the node.

Syntax Description

slot_id	Specifies the slot of the node. A is for the MSM installed in slot A. B is for the MSM installed in slot B.
node_pri	Specifies the priority of the node. The default 0 gives MSM-A a higher priority over MSM-B. The range is 1 to 100; 0 means you have not configured a node priority.

Default

Default node priority is 0.

Usage Guidelines

Use this command to configure the priority of the node. The lower the number, the higher the priority.

The node priority is part of the selection criteria for the primary node. The following list describes the parameters used to determine the primary node:

- Node state—The node state must be STANDBY to participate in leader election and to be selected primary. If the node is in the INIT, DOWN, or FAIL states, the node will not participate in leader election.
- Configuration priority—This is a user assigned priority. The configured priority is compared only after the node meets the minimum thresholds in each category for it to be healthy. Required processes and devices must not fail.
- Software health—This represents the percent of processes available.
- Health of secondary hardware components—This represents the health of switch components, such as the power supplies, fans, and so forth.
- Slot ID—The MSM slot where the node is installed (MSM-A or MSM-B).

If you do not configure any priorities, MSM-A has a higher priority than MSM-B.

Example

The following command configures a priority of 2 for MSM-B:

```
configure node slot B priority 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

configure power supply

```
configure power supply <ps_num> {auto | on}
```

Description

Configures a power supply for either automatic power management, or forced on, regardless of the impact to the total available system power.

Syntax Description

ps_num	Specifies the slot number of the installed power supply unit (PSU) to which this command applies.
auto	Specifies that ExtremeXOS determines the enabled or disabled state of the PSU to maximize total system power. This is the default.
on	Specifies that the PSU be enabled even if ExtremeXOS determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

Default

The default setting is auto; ExtremeXOS either enables or disables the PSU in order to maximize total system power.

Usage Guidelines

If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, ExtremeXOS maximizes system power by automatically taking one of two possible actions:

- If all PSUs are enabled then all PSUs must be budgeted at 110V AC to prevent overload of PSUs with 110V AC inputs.
- OR
- If the PSUs with 110V AC inputs are disabled, then the PSUs with 220V AC inputs can be budgeted with a higher output per PSU.

ExtremeXOS computes the total available power using both methods and automatically uses the PSU configuration that provides the greatest amount of power to the switch. [Table 9](#) lists combinations where ExtremeXOS maximizes system power by disabling the PSUs with 110V AC inputs.

Table 9: PSU combinations where 110V PSUs are disabled

Number of PSUs with 220V AC Inputs	Number of PSUs with 110V AC Inputs
2	1
3	1
3	2
4	1
4	2
5	1

For all other combinations of 220V AC and 110V AC PSUs, ExtremeXOS maximizes system power by enabling all PSUs and budgeting each PSU at 110V AC.

In addition to the PSU, you can specify the following options:

- `auto`—Specifies that ExtremeXOS determines the enabled or disabled state of the PSU to maximize total system power. This is the default.
- `on`—Specifies that the PSU be enabled even if ExtremeXOS determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

You can override automatic power supply management to enable a PSU with 110V AC inputs that ExtremeXOS disables if the need arises, such as for a planned maintenance of 220V AC circuits. If the combination of AC inputs represents one of those listed in [Table 9](#), you can turn on a disabled PSU using the `configure power supply <ps_num> on` command.



NOTE

If you override automatic power supply management, you may reduce the available power and cause one or more I/O modules to power down.

To resume using automatic power supply management on a PSU, use the `configure power supply <ps_num> auto` command. The setting for each PSU is stored as part of the switch configuration.

To display power supply status and power budget information use the `show power` and `show power budget` commands.

Example

The following command configures the PSU in slot 1 to be forced on when either 110V AC or 220V AC power input is present, overriding automatic power management:

```
configure power supply 1 on
```

The switch displays the following message:

```
In a mixed environment of 110V and 220V AC inputs, power management may
automatically disable 110V supplies to maximize the system power budget.
By specifying 'on', you wish to override power management and enable the
specified power supply. This may cause the system power budget to decrease
and one or more I/O cards may be powered off as a result.
```

```
Are you sure you want to continue? (y/n)
```

Enter `y` to continue.

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available only on modular switches.

configure snmp access-profile

```
configure snmp access-profile [<profile_name> | none] {readonly |
readwrite}
```

Description

Configures SNMP to use an ACL policy for access control.

Syntax Description

profile_name	Configures SNMP to use an ACL policy.
none	Cancels a previously configured ACL policy.
readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.

Default

SNMP access is enabled by default, with no ACL policies.

Usage Guidelines

You must be logged in as administrator to configure SNMP parameters.

You can restrict SNMP access by using an ACL and implementing an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for SNMP. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to remove a previously configured ACL.

Creating an ACL Policy File. To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Chapter 11, “Policy Manager,”](#) and [Chapter 12, “Access Lists \(ACLs\),”](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists. To confirm the existence of the policies, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing SNMP Information. To display the current management configuration, including SNMP access related information, whether SNMP access is enabled or disabled, and whether any ACL policies are configured for SNMP, use the following command:

```
show management
```


Example

This example assumes that you already created an ACL to apply to SNMP.

The following command applies the ACL MyAccessProfile_2 to SNMP:

```
configure snmp access-profile MyAccessProfile_2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure snmp add community

```
configure snmp add community [readonly | readwrite] <alphanumeric_string>
```

Description

Adds an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
alphanumeric_string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

Community strings provide a simple method of authentication between a switch and a remote network manager. Read community strings provide read-only access to the switch. The default read-only community string is *public*. Read-write community strings provide read and write access to the switch. The default read/write community string is *private*. Sixteen read-only and sixteen read/write community strings can be configured on the switch, including the defaults.

An authorized trap receiver must be configured to use the correct community strings on the switch for the trap receiver to receive switch-generated traps. In some cases, it may be useful to allow multiple community strings so that all switches and trap receivers are not forced to use identical community strings. The `configure snmp add community` command allows you to add multiple community strings in addition to the default community string.

An SNMP community string can contain up to 32 characters.

Extreme Networks recommends that you change the defaults of the community strings. To change the value of the default read/write and read-only community strings, use the `configure snmp delete community` command.

Example

The following command adds a read/write community string with the value *extreme*:

```
configure snmp add community readwrite extreme
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmp add trapreceiver

```
configure snmp add trapreceiver <ip_address> community [[hex
<hex_community_name>] | <community_name>] {port <port_number>} {from
<src_ip_address>} {mode <trap_mode> [enhanced | standard]}
```

Description

Adds the IP address of a trap receiver to the trap receiver list and specifies which SNMPv1/v2c traps are to be sent.

Syntax Description

ip_address	Specifies an SNMP trap receiver IP address.
hex_community_name	Specifies that the trap receiver is to be supplied as a colon separated string of hex octets.
community_name	Specifies the community string of the trap receiver to be supplied in ASCII format.
port_number	Specifies a UDP port to which the trap should be sent. Default is 162.
src_ip_address	Specifies the IP address of a VLAN to be used as the source address for the trap.
trap_mode	Specifies the mode of the traps: - enhanced—Contains extra varbinds at the end. - standard—Does not contain extra varbinds.

Default

Trap receivers are in enhanced mode by default, and the version is SNMPv2c by default.

Usage Guidelines

The IP address can be unicast, multicast, or broadcast.

An authorized trap receiver can be one or more network management stations on your network. Authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. The switch sends SNMP traps to all trap receivers configured to receive the specific trap group.

To view the SNMP trap receivers configured on the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the destination and community of the SNMP trap receivers configured on the switch.

Example

The following command adds the IP address 10.101.0.100 as a trap receiver with community string *purple*:

```
configure snmp add trapreceiver 10.101.0.100 community purple
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *green*, using port 3003:

```
configure snmp add trapreceiver 10.101.0.105 community green port 3003
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *blue*, and IP address 10.101.0.25 as the source:

```
configure snmp add trapreceiver 10.101.0.105 community blue from 10.101.0.25
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_community_name`, from `<src_ip_address>`, and `{mode <trap_mode> [enhanced | standard]}` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmp delete community

```
configure snmp delete community [readonly | readwrite] [all |  
<alphanumeric_string>]
```

Description

Deletes an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
all	Specifies all of the SNMP community strings.
alphanumeric_string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

You must have at least one community string for SNMP access. If you delete all of the community strings on your system, you will no longer have SNMP access, even if you have SNMP enabled.

The community strings allow a simple method of authentication between the switch and the remote network manager. There are two types of community strings on the switch. Read community strings provide read-only access to the switch. The default read-only community string is *public*. read/write community strings provide read and write access to the switch. The default read/write community string is *private*. Sixteen read-only and sixteen read-write community strings can be configured on the switch, including the defaults. The community string for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. SNMP community strings can contain up to 32 characters.

For increased security, Extreme Networks recommends that you change the defaults of the read/write and read-only community strings.

Use the `configure snmp add` commands to configure an authorized SNMP management station.

Example

The following command deletes a read/write community string named *extreme*:

```
configure snmp delete community readwrite extreme
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmp delete trapreceiver

```
configure snmp delete trapreceiver [{<ip_address> <port_number>}] | {all}]
```

Description

Deletes a specified trap receiver or all authorized trap receivers.

Syntax Description

ip_address	Specifies an SNMP trap receiver IP address.
port_number	Specifies the port associated with the receiver.
all	Specifies all SNMP trap receiver IP addresses.

Default

The default port number is 162.

Usage Guidelines

Use this command to delete a trap receiver of the specified IP address, or all authorized trap receivers.

This command deletes only the first SNMPv1/v2c trap receiver whose IP address and port number match the specified value.

Example

The following command deletes the trap receiver 10.101.0.100 from the trap receiver list:

```
configure snmp delete trapreceiver 10.101.0.100
```

The following command deletes entries in the trap receiver list for 10.101.0.100, port 9990:

```
configure snmp delete trapreceiver 10.101.0.100 9990
```

Any entries for this IP address with a different community string will not be affected.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmp sysContact

```
configure snmp syscontact <sysContact>
```

Description

Configures the name of the system contact.

Syntax Description

sysContact	An alphanumeric string that specifies a system contact name.

Default

N/A.

Usage Guidelines

The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch. A maximum of 255 characters is allowed.

To view the name of the system contact listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system contact.

Example

The following command defines FredJ as the system contact:

```
configure snmp syscontact fredj
```

The following output from the `show switch` command displays *FredJ* as the system contact:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmp sysLocation

```
configure snmp syslocation <sysLocation>
```

Description

Configures the location of the switch.

Syntax Description

sysLocation	An alphanumeric string that specifies the switch location.

Default

N/A.

Usage Guidelines

Use this command to indicate the location of the switch. A maximum of 255 characters is allowed.

To view the location of the switch on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the location of the switch.

Example

The following command configures a switch location name on the system:

```
configure snmp syslocation englab
```

The following output from the `show switch` command displays *englab* as the location of the switch:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmp sysName

```
configure snmp sysname <sysName>
```

Description

Configures the name of the switch.

Syntax Description

sysName	An alphanumeric string that specifies a device name.

Default

The default `sysname` is the model name of the device (for example, `BlackDiamond10808`).

Usage Guidelines

You can use this command to change the name of the switch. A maximum of 32 characters is allowed. The `sysname` appears in the switch prompt. On a SummitStack, the `sysname` appears in the prompt of all active nodes in the stack.

To view the name of the system listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system.

Example

The following command names the switch:

```
configure snmp sysname engineeringlab
```

The following output from the `show switch` command displays *engineeringlab* as the name of the switch:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmpv3 add access

```
configure snmpv3 add access [[hex <hex_group_name>] | <group_name>] {sec-
model [snmpv1 | snmpv2c | usm]} {sec-level [noauth | authnopriv | priv]}
{read-view [[hex <hex_read_view_name>] | <read_view_name>]} {write-view
[[hex <hex_write_view_name>]] | <write_view_name>]} {notify-view [[hex
<hex_notify_view_name>]] | <notify_view_name>]} {volatile}
```

Description

Creates (and modifies) a group and its access rights.

Syntax Description

hex_group_name	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to add or modify. The value is to be supplied in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
read-view	Specifies the read view name: - hex_read_view_name—Specifies a hex value supplied as a colon separated string of hex octets - read_view_name—Specifies an ASCII value
write-view	Specifies the write view name: - hex_write_view_name—Specifies a hex value supplied as a colon separated string of hex octets - write_view_name—Specifies an ASCII value
notify-view	Specifies the notify view name: - hex_notify_view_name—Specifies a hex value supplied as a colon separated string of hex octets - notify_view_name—Specifies an ASCII value
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM
- sec-level—noauth

- read view name—defaultUserView
- write view name— ""
- notify view name—defaultNotifyView
- non-volatile storage

Usage Guidelines

Use this command to configure access rights for a group. All access groups are created with a unique default context, "", as that is the only supported context.

Use more than one character when creating unique community strings and access group names.

A number of default (permanent) groups are already defined. These groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

- The default groups defined (permanent) are *v1v2c_ro* for security name *v1v2c_ro*, *v1v2c_rw* for security name *v1v2c_rw*, *admin* for security name *admin*, and *initial* for security names *initial*, *initialmd5*, *initialsha*, *initialmd5Priv* and *initialshaPriv*.
- The default access defined (permanent) are *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*, and *v1v2cNotifyGroup*.

Example

In the following command, access for the group *defaultROGroup* is created with all the default values: security model *usm*, security level *noauth*, read view *defaultUserView*, no write view, notify view *defaultNotifyView*, and storage *nonvolatile*.

```
configure snmpv3 add access defaultROGroup
```

In the following command, access for the group *defaultROGroup* is created with the values: security model *USM*, security level *authnopriv*, read view *defaultAdminView*, write view *defaultAdminView*, notify view *defaultAdminView*, and storage *nonvolatile*.

```
configure snmpv3 add access defaultROGroup sec-model usm sec-level authnopriv read-  
view defaultAdminView write-view defaultAdminView notify-view defaultAdminView
```

History

This command was first available in ExtremeXOS 10.1.

The *hex_read_view_name*, *hex_write_view_name*, and *hex_notify_view_name* parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add community

```
configure snmpv3 add community [[hex <hex_community_index>] |
<community_index>] name [[hex <hex_community_name>] |<community_name>] user
[[hex <hex_user_name>] | <user_name>] {tag [[hex <hex_transport_tag>] |
<transport_tag>]] {volatile}
```

Description

Adds an SNMPv3 community entry.

Syntax Description

hex_community_index	Specifies the row index in the snmpCommunity table as a hex value supplied as a colon separated string of hex octets.
community_index	Specifies the row index in the snmpCommunity Table as an ASCII value.
hex_community_name	Specifies the community name as a hex value supplied as a colon separated string of hex octets
community_name	Specifies the community name as an ASCII value.
hex_user_name	Specifies the USM user name as a hex value supplied as a colon separated string of hex octets.
user_name	Specifies the USM user name as an ASCII value.
tag	Specifies the tag used to locate transport endpoints in SnmpTargetAddrTable. When this community entry is used to authenticate v1/v2c messages, this tag is used to verify the authenticity of the remote entity. - hex_transport_tag—Specifies a hex value supplied as a colon separated string of hex octets - transport_tag—Specifies an ASCII value
volatile	Specifies volatile storage.

Default

N/A.

Usage Guidelines

Use this command to create or modify an SMMPv3 community in the community MIB.

Example

The following command creates an entry with the community index *comm_index*, community name *comm_public*, and user (security) name *v1v2c_user*:

```
configure snmpv3 add community comm_index name comm_public user v1v2c_user
```

The following command creates an entry with the community index (hex) of *12:0E*, community name (hex) of *EA:12:CD:CF:AB:11:3C*, user (security) name *v1v2c_user*, using transport tag *34872* and volatile storage:

```
configure snmpv3 add community hex 12:0E name hex EA:12:CD:CF:AB:11:3C user v1v2c_user  
tag 34872 volatile
```

History

This command was first available in ExtremeXOS. 10.1.

The `hex_community_index`, `hex_community_name`, `hex_user_name`, and `hex_transport_tag` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add filter

```
configure snmpv3 add filter [[hex <hex_profile_name>] | <profile_name>]
subtree <object_identifier> {/<subtree_mask>} type [included | excluded]
{volatile}
```

Description

Adds a filter to a filter profile.

Syntax Description

hex_profile_name	Specifies the filter profile that the current filter is added to. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile that the current filter is added to in ASCII format.
object identifier	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by <object identifier>/<mask> is to be included.
excluded	Specifies that the MIB subtree defined by <object identifier>/<mask> is to be excluded.
volatile	Specifies volatile storage.

Default

The default values are:

- mask value—empty string (all 1s)
- type—included
- storage—non-volatile

Usage Guidelines

Use this command to create a filter entry in the snmpNotifyFilterTable. Each filter includes or excludes a portion of the MIB. Multiple filter entries comprise a filter profile that can eventually be associated with a target address. Other commands are used to associate a filter profile with a parameter name, and the parameter name with a target address.

This command can be used multiple times to configure the exact filter profile desired.

Example

The following command adds a filter to the filter profile *prof1* that includes the MIB subtree *1.3.6.1.4.1/f0*:

```
configure snmpv3 add filter prof1 subtree 1.3.6.1.4.1/f0 type included
```


History

This command was first available in ExtremeXOS 10.1.

The `hex_profile_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add filter-profile

```
configure snmpv3 add filter-profile [[hex <hex_profile_name>] |
<profile_name>] param [[hex <hex_param_name>]] | <param_name>] {volatile}
```

Description

Associates a filter profile with a parameter name.

Syntax Description

hex_profile_name	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name in ASCII format.
hex_param_name	Specifies a parameter name to associate with the filter profile. The value to follow is to be supplied as a colon separated string of hex octets.
param_name	Specifies a parameter name to associate with the filter profile in ASCII format.
volatile	Specifies volatile storage.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command associates the filter profile *prof1* with the parameter name *P1*:

```
configure snmpv3 add filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_profile_name` and `hex_param_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add group user

```
configure snmpv3 add group [[hex <hex_group_name>] | <group_name>] user
[[hex <hex_user_name>] | <user_name>] {sec-model [snmpv1| snmpv2c | usm]}
{volatile}
```

Description

Adds a user name (security name) to a group.

Syntax Description

hex_group_name	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to add or modify in ASCII format.
hex_user_name	Specifies the user name to add or modify. The value to follow is to be supplies as a colon separated string of hex octets.
user_name	Specifies the user name to add or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM
- non-volatile storage

Usage Guidelines

Use this command to associate a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name *username*, the security name value is the same, *username*.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

Example

The following command associates the user *userV1* to the group *defaultRoGroup* with SNMPv1 security:

```
configure snmpv3 add group defaultRoGroup user userV1 sec-model snmpv1
```

The following command associates the user *userV3* with security model *USM* and storage type *volatile* to the access group *defaultRoGroup*:

```
configure snmpv3 add group defaultRoGroup user userV3 volatile
```

History

This command was first available in ExtremeXOS 10.1.

The *hex_group_name* and *hex_user_name* parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add mib-view

```
configure snmpv3 add mib-view [[hex <hex_view_name>] | <view_name>]
subtree <object_identifier> {/<subtree_mask>} {type [included | excluded]}
{volatile}
```

Description

Adds (and modifies) a MIB view.

Syntax Description

hex_view_name	Specifies the MIB view name to add or modify. The value is to be supplies as a colon separated string of hex octets.
view_name	Specifies the MIB view name to add or modify in ASCII format.
object_identifier	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by <subtree>/<mask> is to be included.
excluded	Specifies that the MIB subtree defined by <subtree>/<mask> is to be excluded.
volatile	Specifies volatile storage.

Default

The default mask value is an empty string (all 1s). The other default values are *included* and *non-volatile*.

Usage Guidelines

Use this command to create a MIB view into a subtree of the MIB. If the view already exists, this command modifies the view to additionally include or exclude the specified subtree.

In addition to the created MIB views, there are three default views. They are of storage type permanent and cannot be deleted, but they can be modified. The default views are: *defaultUserView*, *defaultAdminView*, and *defaultNotifyView*.

Example

The following command creates the MIB view *allMIB* with the subtree *1.3* included as non-volatile:

```
configure snmpv3 add mib-view allMIB subtree 1.3
```

The following command creates the view *extremeMib* with the subtree *1.3.6.1.4.1.1916* included as non-volatile:

```
configure snmpv3 add mib-view extremeMib subtree 1.3.6.1.4.1.1916
```

The following command creates a view *rrpTrapNewMaster* which excludes VRRP notification .1 and the entry is volatile:

```
configure snmpv3 add mib-view rrpTrapNewMaster 1.3.6.1.2.1.68.0.1/ff8 type excluded  
volatile
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add notify

```
configure snmpv3 add notify [[hex <hex_notify_name>] | <notify_name>] tag
[[hex <hex_tag>] | <tag>] {volatile}
```

Description

Adds an entry to the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the notify name to add in ASCII format.
hex_tag	Specifies a string identifier for the notifications to be sent to the target. The value is supplied as a colon separated string of octets.
tag	Specifies a string identifier for the notifications to be sent to the target in ASCII format.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications are sent based on the filters also associated with the target addresses.

Example

The following command sends notifications to addresses associated with the tag *type1*:

```
configure snmpv3 add notify N1 tag type1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name and hex_tag parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add target-addr

```
configure snmpv3 add target-addr [[hex <hex_addr_name> | <addr_name>] param
[[hex <hex_param_name> | <param_name>] ipaddress [[<ip_address>
{<netmask>}] | <ip_address>] {transport-port <port_number> {from
<src_ip_address>} {tag-list <tag_list>} {volatile}
```

Description

Adds and configures an SNMPv3 target address and associates filtering, security, and notifications with that address.

Syntax Description

hex_addr_name	Specifies a string identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address in ASCII format.
hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.
ip_address	Specifies an SNMPv3 target IP address.
port_number	Specifies a UDP port. Default is 162.
src_ip_address	Specifies the IP address of a VLAN to be used as the source address for the trap.
tag-list	Specifies a list of comma separated string identifiers for the notifications to be sent to the target.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- transport-port—port 162
- non-volatile storage

If you do not specify `tag-list` the single tag *defaultNotify*, a pre-defined value in the `snmpNotifyTable`, is used.

Usage Guidelines

Use this command to create an entry in the SNMPv3 `snmpTargetAddressTable`. The `param` parameter associates the target address with an entry in the `snmpTargetParamsTable`, which specifies security and storage parameters for messages to the target address, and an entry in the `snmpNotifyFilterProfileTable`, which specifies filter profiles to use for notifications to the target address. The filter profiles are associated with the filters in the `snmpNotifyFilterTable`.

The list of tag-lists must match one or more of the tags in the `snmpNotifyTable` for the trap to be sent out.

Example

The following command specifies a target address of *10.203.0.22* with the name *A1*, and associates it with the security parameters and target address parameter *P1*:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22
```

The following command specifies a target address of *10.203.0.22* with the name *A1*, and associates it with the security parameters and target address parameter *P1*, and the notification tags *type1* and *type2*:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22 from 10.203.0.23  
tag-list type1,type2
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_addr_name`, `hex_param_name`, `from <src_ip_address>`, and `tag-list <tag_list>` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add target-params

```
configure snmpv3 add target-params [[hex <hex_param_name>] | <param_name>]
user [[hex <hex_user_name>] | <user_name>] mp-model [snmpv1 | snmpv2c |
snmpv3] sec-model [snmpv1 | snmpv2c | usm] {sec-level [noauth | authnopriv
| priv]} {volatile}
```

Description

Adds and configures SNMPv3 target parameters.

Syntax Description

hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.
hex_user_name	Specifies a user name. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies a user name in ASCII format.
mp-model	Specifies a message processing model; choose from SNMPv1, SNMPv2, or SNMPv3.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- sec-level—noauth
- non-volatile storage

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

To associate a target address with a parameter name, see the command `configure snmpv3 add target-addr`.

Example

The following command specifies a target parameters entry named *P1*, a user name of *guest*, message processing and security model of SNMPv2c, and a security level of no authentication:

```
configure snmpv3 add target-params P1 user guest mp-model snmpv2c sec-model snmpv2c  
sec-level noauth
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_param_name` and `hex_user_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add user

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>]
{authentication [md5 | sha] [hex <hex_auth_password> | <auth_password>]}
{privacy [hex <hex_priv_password> | <priv_password>]} {volatile}
```

Description

Adds (and modifies) an SNMPv3 user.

Syntax Description

hex_user_name	Specifies the user name to add or modify. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to add or modify in ASCII format.
MD5	Specifies MD5 authentication.
SHA	Specifies SHA authentication.
authentication	Specifies the authentication password or hex string to use for generating the authentication key for this user.
privacy	Specifies the privacy password or hex string to use for generating the privacy key for this user.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- authentication—no authentication
- privacy—no privacy
- non-volatile storage

Usage Guidelines

Use this command to create or modify an SNMPv3 user configuration.

If hex is specified, supply a 16 octet hex string for MD5, or a 20 octet hex string for SHA.

You must specify authentication if you want to specify privacy. There is no support for privacy without authentication.

The default user names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*. The initial password for *admin* is *password*. For the other default users, the initial password is the user name.

Example

The following command configures the user *guest* on the local SNMP Engine with security level *noauth* (no authentication and no privacy):

```
configure snmpv3 add user guest
```

The following command configures the user *authMD5* to use MD5 authentication with the password *palertyu*:

```
configure snmpv3 add user authMD5 authentication md5 palertyu
```

The following command configures the user *authShapriv* to use SHA authentication with the hex key shown below, the privacy password *palertyu*, and *volatile* storage:

```
configure snmpv3 add user authShapriv authentication sha hex  
01:03:04:05:01:05:02:ff:ef:cd:12:99:34:23:ed:ad:ff:ea:cb:11 privacy palertyu volatile
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add user clone-from

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>] clone-from
[[hex <hex_user_name>] | <user_name>]
```

Description

Creates a new user by cloning from an existing SNMPv3 user.

Syntax Description

hex_user_name	Specifies the user name to add or to clone from. The value is to be supplies as a colon separated string of hex octets.
user_name	Specifies the user name to add or to clone from in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to create a new user by cloning an existing one. After you have successfully cloned the new user, you can modify its parameters using the following command:

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>] {authentication [md5 |
sha] [hex <hex_auth_password> | <auth_password>]} {privacy [hex <hex_priv_password> |
<priv_password>]} {volatile}
```

Users cloned from the default users will have the storage type of non-volatile. The default names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Example

The following command creates a user *cloneMD5* with same properties as the default user *initialmd5*. All authorization and privacy keys will initially be the same as with the default user *initialmd5*.

```
configure snmpv3 add user cloneMD5 clone-from initialmd5
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete access

```
configure snmpv3 delete access [all-non-defaults | {[hex <hex_group_name>]
| <group_name>] {sec-model [snmpv1 | snmpv2c | usm] sec-level [noauth |
authnopriv | priv]]}]
```

Description

Deletes access rights for a group.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) security groups are to be deleted.
hex_group_name	Specifies the group name to be deleted. The value is to be supplies as a colon separated string of hex octets.
group_name	Specifies the group name to be deleted in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.

Default

The default values are:

- sec-model—USM
- sec-level—noauth

Usage Guidelines

Use this command to remove access rights for a group. Use the `all-non-defaults` keyword to delete all the security groups, except for the default groups. The default groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

Deleting an access will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[hex <hex_group_name>] | <group_name>}} user [all-non-
defaults | {[hex <hex_user_name>] | <user_name>] {sec-model [snmpv1|snmpv2c|usm]]}]
```

Example

The following command deletes all entries with the group name *userGroup*:

```
configure snmpv3 delete access userGroup
```

The following command deletes the group *userGroup* with the security model *snmpv1* and security level of authentication and no privacy (*authnopriv*):

```
configure snmpv3 delete access userGroup sec-model snmpv1 sec-level authnopriv
```

History

This command was first available in ExtremeXOS 10.1.

The *hex_group_name* parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete community

```
configure snmpv3 delete community [all-non-defaults | {[[hex
<hex_community_index>] | <community_index>} | {name [[hex
<hex_community_name>] | <community_name>}]]
```

Description

Deletes an SNMPv3 community entry.

Syntax Description

all-non-defaults	Specifies that all non-default community entries are to be removed.
hex_community_index	Specifies the row index in the snmpCommunityTable. The value is to be supplied as a colon separated string of hex octets.
community_index	Specifies the row index in the snmpCommunityTable in ASCII format.
hex_community_name	Specifies the community name. The value is to be supplied as a colon separated string of hex octets.
community_name	Specifies the community name in ASCII format.

Default

The default entries are *public* and *private*.

Usage Guidelines

Use this command to delete an SMMPv3 community in the community MIB.

Example

The following command deletes an entry with the community index *comm_index*:

```
configure snmpv3 delete community comm_index
```

The following command creates an entry with the community name (hex) of *EA:12:CD:CF:AB:11:3C*:

```
configure snmpv3 delete community name hex EA:12:CD:CF:AB:11:3C
```

History

This command was first available in ExtremeXOS 10.1.

The *hex_community_index* and *hex_community_name* parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete filter

```
configure snmpv3 delete filter [all | [[hex <hex_profile_name>] |  
<profile_name>] {subtree <object_identifier>}]]
```

Description

Deletes a filter from a filter profile.

Syntax Description

all	Specifies all filters.
hex_profile_name	Specifies the filter profile of the filter to delete. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile of the filter to delete in ASCII format.
object_identifier	Specifies the MIB subtree of the filter to delete.

Default

N/A.

Usage Guidelines

Use this command to delete a filter entry from the snmpNotifyFilterTable. Specify `all` to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a subtree to delete just those entries for that filter profile and subtree.

Example

The following command deletes the filters from the filter profile *prof1* that reference the MIB subtree *1.3.6.1.4.1*:

```
configure snmpv3 delete filter prof1 subtree 1.3.6.1.4.1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_profile_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete filter-profile

```
configure snmpv3 delete filter-profile [all | [[hex <hex_profile_name>] |
<profile_name>] {param [[hex <hex_param_name>] | <param_name>}]}
```

Description

Removes the association of a filter profile with a parameter name.

Syntax Description

all	Specifies all filter profiles.
hex_profile_name	Specifies the filter profile name to delete. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name to delete in ASCII format.
hex_param_name	Specifies to delete the filter profile with the specified profile name and parameter name. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies to delete the filter profile with the specified profile name and parameter name in ASCII format.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to delete entries from the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. Specify `all` to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a parameter name to delete just those entries for that filter profile and parameter name.

Example

The following command deletes the filter profile *prof1* with the parameter name *P1*:

```
configure snmpv3 delete filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_profile_name` and `hex_param_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete group user

```
configure snmpv3 delete group {[[hex <hex_group_name>] | <group_name>]}
user [all-non-defaults | {[[hex <hex_user_name>] | <user_name>] {sec-model
[snmpv1|snmpv2c|usm]}]}
```

Description

Deletes a user name (security name) from a group.

Syntax Description

hex_group_name	Specifies the group name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to delete or modify in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted from the group.
hex_user_name	Specifies the user name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to delete or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).

Default

The default value for sec-model is USM.

Usage Guidelines

Use this command to remove the associate of a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name *username*, the security name value is the same, *username*.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

The default groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

The default users are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Example

The following command deletes the user *guest* from the group *UserGroup* for the security model *snmpv2c*:

```
configure snmpv3 delete group UserGroup user guest sec-model snmpv2c
```

The following command deletes the user *guest* from the group *userGroup* with the security model *USM*:

```
configure snmpv3 delete group userGroup user guest
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_group_name` and the `hex_user_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete mib-view

```
configure snmpv3 delete mib-view [all-non-defaults | {[[hex
<hex_view_name>] | <view_name>] {subtree <object_identifier>}}]
```

Description

Deletes a MIB view.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) MIB views are to be deleted.
hex_view_name	Specifies the MIB view to delete. The value is to be supplied as a colon separated string of hex octets.
view_name	Specifies the MIB view name to delete in ASCII format.
object_identifier	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to delete a MIB view. Views which are being used by security groups cannot be deleted. Use the `all-non-defaults` keyword to delete all the MIB views (not being used by security groups) except for the default views. The default views are: *defaultUserView*, *defaultAdminView*, and *defaultNotifyView*.

Use the `configure snmpv3 add mib-view` command to remove a MIB view from its security group, by specifying a different view.

Example

The following command deletes all views (only the permanent views will not be deleted):

```
configure snmpv3 delete mib-view all-non-defaults
```

The following command deletes all subtrees with the view name *AdminView*:

```
configure snmpv3 delete mib-view AdminView
```

The following command deletes the view *AdminView* with subtree 1.3.6.1.2.1.2

```
configure snmpv3 delete mib-view AdminView subtree 1.3.6.1.2.1.2
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete notify

```
configure snmpv3 delete notify [{"hex <hex_notify_name>" | <notify_name>"}  
| all-non-defaults]
```

Description

Deletes an entry from the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the notify name to add in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) notifications are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications will be sent, based on the filters also associated with the target addresses.

The one default notification that cannot be deleted is *defaultNotify*.

Example

The following command removes the *N1* entry from the table:

```
configure snmpv3 delete notify N1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_notify_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete target-addr

```
configure snmpv3 delete target-addr [{{[[hex <hex_addr_name>] |  
<addr_name>]]} | all]
```

Description

Deletes SNMPv3 target addresses.

Syntax Description

hex_addr_name	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address.
all	Specifies all target addresses.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetAddressTable.

Example

The following command deletes target address named *A1*:

```
configure snmpv3 delete target-addr A1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_addr_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete target-params

```
configure snmpv3 delete target-params [{{[[hex <hex_param_name>] |  
<param_name>]]} | all]
```

Description

Deletes SNMPv3 target parameters.

Syntax Description

hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

Example

The following command deletes a target parameters entry named *P1*:

```
configure snmpv3 delete target-params P1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_param_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete user

```
configure snmpv3 delete user [all-non-defaults | [[hex <hex_user_name>] |
<user_name>]]
```

Description

Deletes an existing SNMPv3 user.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted.
hex_user_name	Specifies the user name to delete. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to delete.

Default

N/A.

Usage Guidelines

Use this command to delete an existing user.

Use the `all-non-defaults` keyword to delete all users, except for the default (permanent) users. The default user names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Deleting a user will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group [[hex <hex_group_name>] | <group_name>]] user [all-non-
defaults | [[hex <hex_user_name>] | <user_name>] {sec-model [snmpv1|snmpv2c|usm]}]]
```

Example

The following command deletes all non-default users:

```
configure snmpv3 delete user all-non-defaults
```

The following command deletes the user *guest*:

```
configure snmpv3 delete user guest
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 engine-boots

```
configure snmpv3 engine-boots <(1-2147483647)>
```

Description

Configures the SNMPv3 Engine Boots value.

Syntax Description

(1-2147483647)	Specifies the value of engine boots.

Default

N/A.

Usage Guidelines

Use this command if the Engine Boots value needs to be explicitly configured. Engine Boots and Engine Time will be reset to zero if the Engine ID is changed. Engine Boots can be set to any desired value but will latch on its maximum, 2147483647.

Example

The following command configures Engine Boots to 4096:

```
configure snmpv3 engine-boots 4096
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmpv3 engine-id

```
configure snmpv3 engine-id <hex_engine_id>
```

Description

Configures the SNMPv3 `snmpEngineID`.

Syntax Description

<code>hex_engine_id</code>	Specifies the colon delimited hex octet that serves as part of the <code>snmpEngineID</code> (5-32 octets).

Default

The default `snmpEngineID` is the device MAC address.

Usage Guidelines

Use this command if the `snmpEngineID` needs to be explicitly configured. The first four octets of the ID are fixed to 80:00:07:7C, which represents Extreme Networks Vendor ID. Once the `snmpEngineID` is changed, default users will be reverted back to their original passwords/keys, while non-default users will be reset to the security level of no authorization, no privacy.

In a chassis, the `snmpEngineID` will be generated using the MAC address of the MSM with which the switch boots first. For MSM hitless failover, the same `snmpEngineID` will be propagated to both of the MSMs.

Example

The following command configures the `snmpEngineID` to be 80:00:07:7C:00:0a:1c:3e:11:

```
configure snmpv3 engine-id 00:0a:1c:3e:11
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure sntp-client

```
configure sntp-client [primary | secondary] <host-name-or-ip> {vr
<vr_name>}
```

Description

Configures an NTP server for the switch to obtain time information.

Syntax Description

primary	Specifies a primary server name.
secondary	Specifies a secondary server name.
host-name-or-ip	Specifies a host name or IP address.
vr	Specifies use of a virtual router. NOTE: The BlackDiamond 8800 series switch (formerly known as Aspen), SummitStack, and the Summit family of switches do not support user-created VRs.
vr_name	Specifies the name of a virtual router.

Default

N/A.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Queries are first sent to the primary server. If the primary server does not respond within 1 second, or if it is not synchronized, the switch queries the second server. If the switch cannot obtain the time, it restarts the query process. Otherwise, the switch waits for the `sntp-client update interval` before querying again.

Example

The following command configures a primary NTP server:

```
configure sntp-client primary 10.1.2.2
```

The following command configures the primary NTP server to use the management virtual router VR-Mgmt

```
configure sntp-client primary 10.1.2.2 vr VR-Mgmt
```

History

This command was first available in ExtremeXOS 10.1.

The `vr <vr_name>` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sntp-client update-interval

```
configure sntp-client update-interval <update-interval>
```

Description

Configures the interval between polls for time information from SNTP servers.

Syntax Description

update-interval	Specifies an interval in seconds.

Default

64 seconds.

Usage Guidelines

None.

Example

The following command configures the interval timer:

```
configure sntp-client update-interval 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure telnet access-profile

```
configure telnet access-profile [<access_profile> | none]
```

Description

Configures Telnet to use an ACL policy for access control.

Syntax Description

access_profile	Specifies an ACL policy.
none	Cancels a previously configured ACL policy.

Default

Telnet is enabled with no ACL policies and uses TCP port 23.

Usage Guidelines

You must be logged in as administrator to configure Telnet parameters.

You can restrict Telnet access by using an ACL and implementing an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the Telnet port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to remove a previously configured ACL.

Creating an ACL Policy File. To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Chapter 11, “Policy Manager,”](#) and [Chapter 12, “Access Lists \(ACLs\),”](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing Telnet Information. To display the status of Telnet, including the current TCP port, the virtual router used to establish a Telnet session, and whether ACLs are controlling Telnet access, use the following command:

```
show management
```

Example

This example assumes that you already created an ACL to apply to Telnet.

The following command applies the ACL MyAccessProfile_2 to Telnet:

```
configure telnet access-profile MyAccessProfile_2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure telnet port

```
configure telnet port [<portno> | default]
```

Description

Configures the TCP port used by Telnet for communication.

Syntax Description

portno	Specifies a TCP port number. The default is 23. The range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023.
default	Specifies the default Telnet TCP port number. The default is 23.

Default

The switch listens for Telnet connections on Port 23.

Usage Guidelines

You must be logged in as administrator to configure the Telnet port.

The `portno` range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023. If you attempt to configure a reserved port, the switch displays an error message similar to the following:

```
configure telnet port 22
Error: port number is a reserved port
```

If this occurs, select a port number that is not a reserved port.

Beginning with ExtremeXOS 11.2, the switch accepts IPv6 connections.

Example

The following command changes the port used for Telnet to port 85:

```
configure telnet port 85
```

The following command returns the port used for Telnet to the default port of 23:

```
configure telnet port default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure telnet vr

```
configure telnet vr [all | default | <vr_name>]
```

Description

Configures the virtual router used on the switch for listening for Telnet connections.

Syntax Description

all	Specifies to use all virtual routers for Telnet connections.
default	Specifies to use the default virtual router for Telnet connections. The default router is VR-Mgmt.
vr_name	Specifies the name of the virtual router to use for Telnet connections. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The default is all vr's.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

You must be logged in as administrator to configure the virtual router.

Beginning with ExtremeXOS 11.2, the switch accepts IPv6 connections.

If you specify `all`, the switch listens on all of the available virtual routers for Telnet connections.

The `vr_name` specifies the name of the virtual router to use for Telnet connections.

If you specify a virtual router name that does not exist, the switch displays an error message similar to the following:

```
configure telnet vr vr-ttt
                        ^
%% Invalid input detected at '^' marker.
```

Example

The following command configures the switch to listen for and receive Telnet requests on all virtual routers:

```
configure telnet vr all
```

History

This command was first available in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable dhcp vlan

```
disable dhcp vlan [<vlan_name> | all]
```

Description

Disables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs

Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command disables the generation and processing of DHCP packets on a VLAN named *accounting*:

```
disable dhcp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable snmp access

```
disable snmp access {snmp-v1v2c}
```

Description

Selectively disables SNMP on the switch.

Syntax Description

snmp-v1v2c	Disables SNMPv1/v2c access only; does not affect SNMPv3 access.

Default

Enabled.

Usage Guidelines

Disabling SNMP access does not affect the SNMP configuration (for example, community strings). However, if you disable SNMP access, you will be unable to access the switch using SNMP.

To allow access, use the following command:

```
enable snmp access
```

By using the enable and disable commands you can enable all SNMP access, no SNMP access, or only SNMPv3 access. You cannot enable only SNMPv1/v2c access. To enable SNMPv3 only access on the switch, use the following commands:

```
enable snmp access  
disable snmp access snmp-v1v2c
```

Example

The following command disables all SNMP access on the switch:

```
disable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable snmp traps

```
disable snmp traps
```

Description

Prevents SNMP traps from being sent from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command does not clear the SNMP trap receivers that have been configured. The command prevents SNMP traps from being sent from the switch even if trap receivers are configured.

To view if SNMP traps are being sent from the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command prevents SNMP traps from being sent from the switch to the trap receivers:

```
disable snmp traps
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable sntp-client

```
disable sntp-client
```

Description

Disables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command disables the SNTP client:

```
disable sntp-client
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable telnet

```
disable telnet
```

Description

Disables external Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.



NOTE

Telnet sessions between MSMs or the nodes of a stack are not affected by this command.

Example

With administrator privilege, the following command disables external Telnet services on the switch:

```
disable telnet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable watchdog

```
disable watchdog
```

Description

Disables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This can be caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command disables the watchdog timer:

```
disable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable dhcp vlan

```
enable dhcp vlan [<vlan_name> | all]
```

Description

Enables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command enables the generation and processing of DHCP packets on a VLAN named *accounting*:

```
enable dhcp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable snmp access

```
enable snmp access
```

Description

Turns on SNMP support for SNMPv3 and v1/v2c on the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

To have access to the SNMP agent residing in the switch, at least one VLAN must have an IP address assigned to it.

Any network manager running SNMP can manage the switch (for v1/v2c), provided the MIB is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

For SNMPv3, additional security keys are used to control access, so an SNMPv3 manager is required for this type of access.

This command enables both v1/v2c and v3 access, so the switch can be accessed with either method. Use the following commands to allow only v3 access:

```
enable snmp access
disable snmp access snmp-v1v2c
```

Use the following command to prevent any SNMP access:

```
disable snmp access
```

There is no way to disable v3 access and allow v1/v2c access.

ExtremeXOS 11.2 introduces the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for SNMP—the default setting is enabled—the switch returns the following interactive script:

```
Since you have chosen less secure management methods, please remember to increase
the security of your network by taking the following actions:
```

- * change your admin password
- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

In addition, you can return to safe defaults mode by issuing the following commands:

- `synchronize stacking all`
- `configure safe-default-script`

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see [“Safe Defaults Setup Method” on page 53](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command enables all SNMP access for the switch:

```
enable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable snmp traps

```
enable snmp traps
```

Description

Turns on SNMP trap support.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers.

To view if SNMP traps are being sent from the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command enables SNMP trap support on the switch:

```
enable snmp traps
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable sntp-client

```
enable sntp-client
```

Description

Enables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command enables the SNTP client:

```
enable sntp-client
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable telnet

```
enable telnet
```

Description

Enables external Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.

ExtremeXOS 11.2 introduces the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for Telnet—the default setting is enabled—the switch returns the following interactive script:

Since you have chosen less secure management methods, please remember to increase the security of your network by taking the following actions:

- * change your admin password
- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

In addition, you can return to safe defaults mode by issuing the following commands:

- `synchronize stacking all`
- `configure safe-default-script`

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see [“Safe Defaults Setup Method” on page 53](#) in the *ExtremeXOS Concepts Guide*.

Example

With administrator privilege, the following command enables Telnet services on the switch:

```
enable telnet
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable watchdog

```
enable watchdog
```

Description

Enables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This is caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command enables the watchdog timer:

```
enable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

exit

exit

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
exit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

logout

logout

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
logout
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

quit

```
quit
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
quit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show dhcp-client state

```
show dhcp-client state
```

Description

Displays the current DHCP/BOOTP client state for each vlan.

Syntax Description

This command has no arguments or variables.

Default

Displays the client state for all existing VLANs.

Usage Guidelines

None.

Example

The following command displays the DHCP/BOOTP status for all VLANs:

```
show dhcp-client state
```

Depending on your configurations, output from this command is similar to the following:

Client VLAN	Protocol	Server	Current State
Default	BOOTP	10.1.2.3	Received IP address configured on vlan
accounting	DHCP	10.2.3.4	DHCP state; Requesting
Mgmt	None	0.0.0.0	

A total of 3 vlan(s) where displayed

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show checkpoint-data

```
show checkpoint-data {<process>}
```

Description

Displays the status of one or more processes being copied from the primary MSM to the backup MSM.

Syntax Description

process	Specifies the name of the processes being copied.

Default

N/A.

Usage Guidelines

This command displays, in percentages, the amount of internal state copying completed by each process and the traffic statistics between the process on both the primary and the backup MSMs.

This command is also helpful in debugging synchronization problems that occur at run-time. To check the status of synchronizing the MSMs, use the `show switch` command.

Depending on the software version running on your switch and the type of switch you have, additional or different checkpoint status information may be displayed.

Example

The following command displays the checkpointing status and the traffic statics of all of the processes between the primary and the backup MSM:

```
show checkpoint-data
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	% Chkpt	Debug-info
devmgr	3812	1731	0	3	3	100% ON OK	1 (00008853)
dirser	0	0	0	0	0	0% ON OK	1 (000008D3)
ems	5	0	0	0	0	100% ON OK	1 (000008D3)
nodemgr	0	0	0	0	0	0% ON OK	1 (000008D3)
snmpSubagent	0	0	0	0	0	0% ON OK	1 (000018D3)
snmpMaster	0	0	0	0	0	0% ON OK	1 (000008D3)
cli	0	0	0	0	0	0% ON OK	1 (000018D3)
edp	0	0	0	0	0	0% ON OK	1 (000008D3)
cfgmgr	82	82	0	1	1	100% ON OK	1 (000018D3)
elrp	0	0	0	0	0	0% ON OK	1 (000008D3)
vlan	1047	1	0	0	0	100% ON OK	1 (000008D3)
aaa	0	0	0	0	0	0% ON OK	1 (000008D3)

```

fdb          957      2      0      0      0 100% ON OK 1 (000008D3)
msgsrv       0       0      0      0      0 100% ON OK 1 (000008D3)
eaps         0       0      0      0      0  0% ON OK 1 (000008D3)
stp          1       0      0      0      0  0% ON OK 1 (000008D3)
esrp         1       0      0      0      0 100% ON OK 1 (000008D3)
polMgr       0       0      0      0      0  0% ON OK 1 (000008D3)
mcmgr        2       2      0      0      0 100% ON OK 1 (000008D3)
acl          0       0      0      0      0 100% ON OK 1 (000008D3)
netLogin     0       0      0      0      0  0% ON OK 1 (000008D3)
ospf         0       0      0      0      0  0% ON OK 1 (000008D3)
netTools     1       0      0      0      0 100% ON OK 1 (000008D3)
telnetd      0       0      0      0      0  0% ON OK 1 (000008D3)
rtmgr        4       4      0      0      0 100% ON OK 1 (000008D3)
vrrp        378      0      0      0      0  0% ON OK 1 (000008D3)
tftpd        0       0      0      0      0  0% ON OK 1 (000008D3)
thttpd       0       0      0      0      0  0% ON OK 1 (000008D3)
rip          0       0      0      0      0  0% ON OK 1 (000008D3)
dosprotect   0       0      0      0      0  0% ON OK 1 (000008D3)
epm          0       0      0      0      0  0% ON OK 1 (000008D3)
hal          0       0      0      0      0  0% ON OK 1 (000008D3)
bgp          0       0      0      0      0  0% ON OK 1 (000008D3)
pim          0       0      0      0      0  0% ON OK 1 (000008D3)
etmon       185     185      0      0      0 100% ON OK 1 (000008D3)

```

To view the output for a specific process, use the `process` option. The following command displays detailed information for the STP process:

```
show checkpoint-data stp
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	% Chkpt	Debug-info
stp	1	0	0	0	0	0% ON OK	1 (000008D3)

History

This command was first available in ExtremeXOS 10.1.

An error count was added to the output in ExtremeXOS 11.1.

Platform Availability

This command is available only on modular switches.

show management

```
show management
```

Description

Displays the SNMP settings configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

The following management output is displayed:

- Enable/disable state for Telnet, and SNMP access
- Login statistics
 - Enable/disable state for idle timeouts
 - Maximum number of CLI sessions
- SNMP community strings
- SNMP trap receiver list

For ExtremeXOS 11.0 and later, the following management output is also displayed:

- SNMP trap receiver source IP address
- SNMP statistics counter
- SSH access states of enabled, disabled, and module not loaded
- CLI configuration logging
- SNMP access states of v1, v2c disabled and v3 enabled

If all three types of SNMP access are enabled or disabled, SNMP access is displayed as either Enabled or Disabled.

For ExtremeXOS 11.1 and later, the following management output is also displayed:

- Enable/disable state for RMON

For ExtremeXOS 11.2 and later, the following management output is also displayed:

- Access-profile usage configured via Access Control Lists (ACLs) for additional Telnet and SSH2 security

For ExtremeXOS 11.6 and later, the following management output is also displayed:

- CLI scripting settings
 - Enable/disable state
 - Error message setting
 - Persistence mode

Example

The following command displays configured SNMP settings on the switch:

```
show management
```

The following is sample output from this command:

```
CLI idle timeout           : Disabled
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI scripting              : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
Telnet access              : Enabled (tcp port 23 vr all)
                           : Access Profile : not set

SSH access                 : Enabled (Key valid, tcp port 22 vr all)
                           : Access Profile : not set

SNMP access                : Enabled
Total Read Only Communities : 1
Total Read Write Communities : 1
RMON                       : Disabled
SNMP Traps                 : Enabled
SNMP v1/v2c TrapReceivers  :
    Destination      Source IP Address      Flags
    192.168.1.102 /10550          2E
    10.255.48.38 /10550          2E
    10.255.43.12 /10550          2E
    10.209.10.202 /10550          2E

Flags:  Version: 1=v1 2=v2c
        Mode: S=Standard E=Enhanced

SNMP stats:      InPkts 231      OutPkts 248      Errors 4      AuthErrors 0
                  Gets   12      GetNexts 218     Sets   1
SNMP traps:      Sent   17      AuthTraps Enabled
```

History

This command was first available in ExtremeXOS 10.1.

The trap receiver source IP address, SNMP counter statistics, SSH access, CLI logging, and SNMP access states was added to the output in ExtremeXOS 11.0.

The enabled/disable state for RMON was added to the output in ExtremeXOS 11.1.

Additional Telnet and SSH2 information about ACL usage was added to the output in ExtremeXOS 11.2.

Information about CLI scripting including, the enabled/disabled state, error mode, and persistent mode was added to the output in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show node

```
show node {detail}
```

Description

Displays the status of the nodes in the system as well as the general health of the system.

Syntax Description

detail	Displays the information on a per-node basis rather than in a tabular format.

Default

N/A.

Usage Guidelines

Use this command to display the current status of the nodes and the health of the system. The information displayed shows the node configurations (such as node priority) and the system and hardware health computations. You can use this information to determine which node will be elected primary in case of a failover.

[Table 10](#) lists the node statistic information collected by the switch.

Table 10: Node states

Node State	Description
BACKUP	In the backup state, this node becomes the primary node if the primary fails or enters the DOWN state. The backup node also receives the checkpoint state data from the primary.
DOWN	In the down state, the node is not available to participate in leader election. The node enters this state during any user action, other than a failure, that makes the node unavailable for management. Examples of user actions are: - Upgrading the software - Rebooting the system using the <code>reboot</code> command - Initiating an MSM failover using the <code>run msm-failover</code> command - Synchronizing the MSMs software and configuration in non-volatile storage using the <code>synchronize</code> command
FAIL	In the fail state, the node has failed and needs to be restarted or repaired. The node reaches this state if the system has a hardware or software failure.
INIT	In the initial state, the node is being initialized. A node stays in this state when it is coming up and remains in this state until it has been fully initialized. Being fully initialized means that all of the hardware has been initialized correctly and there are no diagnostic faults.
MASTER	In the primary state, the node is responsible for all switch management functions.
STANDBY	In the standby state, leader election occurs—the primary and backup nodes are elected. The priority of the node is only significant in the standby state.

Example

The following command displays the status of the node, the priority of the node, and the general health of the system:

```
show node
```

The following is sample output from this command:

Node	State	Priority	SwHealth	HwHealth
MSM-A	MASTER	0	49	7
MSM-B	BACKUP	0	49	7

If you specify the `detail` option, the same information is displayed on a per node basis rather than in a tabular format.

Node MSM-A information:

```
Node State:    MASTER
Node Priority:  0
Sw Health:     49
Hw Health:     7
```

Node MSM-B information:

```
Node State:    BACKUP
Node Priority:  0
Sw Health:     49
Hw Health:     7
```

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

show odometers

show odometers

Description

Displays a counter for each component of a switch that shows how long it has been functioning since it was manufactured.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output from this command displays how long individual components in the switch have been functioning since it was manufactured. This odometer counter is kept in the EEPROM of each monitored component. On a modular switch, this means that even if you plug in the component into a different chassis, the odometer counter is available in the new switch chassis.

Monitored Components.

On a modular switch, the odometer monitors the following components:

- Chassis
- MSMs
- I/O modules
- Power controllers

On the Summit family switches (whether or not included in a SummitStack), the odometer monitors the following components:

- Switch
- XGM-2xn card

Recorded Statistics. The following odometer statistics are collected by the switch:

- Service Days—The amount of days that the component has been running
- First Recorded Start Date—The date that the component was powered-up and began running

Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different odometer information may be displayed.

Example

The following command displays how long each component of a switch has been functioning since its manufacture date:

```
show odometers
```

The following is sample output from a BlackDiamond 10808 switch:

Field Replaceable Units	Service Days	First Recorded Start Date
Chassis : BD-10808	107	Feb-23-2004
Slot-1 : G60X	99	Dec-10-2003
Slot-2 : G60X	74	Mar-22-2004
Slot-3 : G60X	151	Jan-12-2004
Slot-4 :		
Slot-5 : 10G6X	49	Apr-09-2004
Slot-6 :		
Slot-7 : G60T	184	Dec-03-2003
Slot-8 : 10G6X	146	Jan-12-2004
MSM-A : MSM-1XL	62	Apr-21-2004
MSM-B : MSM-1XL	172	Dec-14-2003
PSUCTRL-1 :	152	Mar-17-2004
PSUCTRL-2 :		

The following is sample output from the BlackDiamond 8800 series switch:

Field Replaceable Units	Service Days	First Recorded Start Date
Chassis : BD-8810	209	Dec-07-2004
Slot-1 : G48T	208	Dec-07-2004
Slot-2 : 10G4X	219	Nov-02-2004
Slot-3 : G48T	228	Oct-26-2004
Slot-4 : G24X	226	Oct-19-2004
Slot-5 : G8X	139	Dec-07-2004
Slot-6 :		
Slot-7 : 10G4X	160	Dec-16-2004
Slot-8 : 10G4X	133	Dec-14-2004
Slot-9 : G48P	111	Nov-04-2004
Slot-10 :		
MSM-A : MSM-G8X	137	Dec-07-2004
MSM-B :		
PSUCTRL-1 :	209	Dec-07-2004
PSUCTRL-2 :	208	Dec-07-2004

The following is sample output from a stand-alone Summit series switch:

Field Replaceable Units	Service Days	First Recorded Start Date
Switch : SummitX450-24t	7	Dec-08-2004
XGM-2xn-1 :		

History

This command was first available in ExtremeXOS 10.1.

Information about the power controller(s) for modular switches was added to the `show odometers` output in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show power

```
show power {<ps_num>} {detail}
```

Description

Displays the current status of the installed power supplies.

Command Syntax

ps_num	Specifies the slot number of the installed power supply.
detail	The detail option is reserved for future use.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supplies.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects the following power supply information:

- State—Indicates the current state of the power supply. Options are:
 - Empty—There is no power supply installed.
 - Power Failed—The power supply has failed.
 - Powered Off—The power supply is off.
 - Powered On—The power supply is on and working normally.

Modular switches only:

Located next to the “State” of the power supply, the following information provides more detailed status information. Options are:

- Disabled for net power gain—Indicates that the power supply is disabled in order to maximize the total available system power
- Configured ON—Indicates that the user requested to enable a disabled power supply regardless of the affect on the total available system power
- Configured ON when present—Indicates that the power supply slot is currently empty, but the user requested to enable the power supply regardless of the affect on the total available system power
- Unsupported—Indicates that a 600/900 W AC PSU is inserted in a chassis other than the BlackDiamond 8806.

- PartInfo—Provides information about the power supply. Depending on your switch, options include:

Modular switches only:

- Serial number—A collection of numbers and letters, that make up the serial number of the power supply.
- Part number—A collection of numbers and letters that make up the part number of the power supply.

Summit family switches only:

- Internal Power Supply (PowerSupply 1 information)—The Summit family switches come with one power supply pre-installed at the factory. The Summit power supply is not user-replaceable; therefore, the part information display indicates internal power supply.
- External Power Supply (PowerSupply 2-4 information)—Displays information about the optional External Power System (EPS) that allows you to add a redundant power supply to the Summit family switches to protect against a power supply failure.

On modular switches, the output also includes the following information:

- Revision—Displays the revision number of the power supply.
- Odometer—Specifies the date and how long the power supply has been operating.
- Temperature—Specifies, in celsius, the current temperature of the power supply.
- Input—Specifies the input voltage and the current requirements of the power supply and whether the input is AC or DC.
- Output 1 and Output 2—Specifies the output voltage and the current supplied by the power supply. The values are only displayed if known for the platform.

SummitStack displays the supplies associated to each active node that is present in a slot. The supplies are represented with flags that describe whether the supply is providing power, has failed, or is providing no power, or if the supply has had its 48v power output automatically turned off because two or three external power supplies are available. For more information, see the [“show power \(SummitStack Only\)”](#) section in [Chapter 16, “SummitStack Commands.”](#)

In ExtremeXOS 10.1 and earlier, use the `show powersupplies {detail}` command to view detailed health information about the power supplies.

Example

Modular switch example:

The following command displays the status of the power supply installed in slot 1 in a modular switch:

```
show power 1
```

The following is sample output from this command:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        PS 2336 5003J-00479 4300-00137
Revision:        2.0
Odometer:        90 days 5 hours
Temperature:     29.0 deg C
Fan 1:           6473 RPM
Fan 2:           6233 RPM
Input:           230.00 V AC
```

```
Output 1:      48.50 V,  7.25 A   (48V/1104W Max)
Output 2:      12.44 V,  0.62 A   (12V/48W Max)
```

If power management needs to disable a power supply to maximize the total available power, you see Disabled for net power gain next to the state of the power supply, as shown in the sample truncated output:

```
PowerSupply 1 information:
State:      Powered Off (Disabled for net power gain)
PartInfo:   PS 2336 0413J-00732 4300-00137
...
```

If you choose to always enable a power supply, regardless of the affect on the total available power, you see Configured ON next to the state of the power supply, as shown in the sample truncated output:

```
PowerSupply 1 information:
State:      Powered On  (Configured ON)
PartInfo:   PS 2336 0413J-00732 4300-00137
```

If you install the 600/900 W AC PSU in a chassis other than a BlackDiamond 8806, you see unsupported next to the state of the power supply, as shown in this sample truncated output:

```
PowerSupply 3 information:
State:      Unsupported
PartInfo:   PS 2431 0622J-00013 4300-00161
```

Summit family switches example:

The following command displays the status of the power supplies installed in the Summit X450 series switch:

```
show power
```

The following sample output assumes that you have not installed an EPS:

```
PowerSupply 1 information:
State:      Powered On
PartInfo:   Internal Power Supply
PowerSupply 2 information:
State:      Empty
```

The following sample output assumes that you have installed an EPS:

```
PowerSupply 1 information:
State:      Powered On
PartInfo:   Internal Power Supply

PowerSupply 2 information:
State:      Powered On
PartInfo:   External Power Supply
```

The following sample output for the Summit X450e-48p assumes a connection to an external EPS-C with three EPS-600LS modules installed:

```
show power
```

```
PowerSupply 1 information:
```

```
State:      Powered Off
PartInfo:   Internal Power Supply
Input:      0.00 V AC
```

```
PowerSupply 2 information:
```

```
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC
```

```
PowerSupply 3 information:
```

```
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC
```

```
PowerSupply 4 information:
```

```
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC
```

History

This command was first available in an ExtremeXOS 10.1.

The syntax for this command was modified in ExtremeXOS 11.0 from `show powersupplies` to `show power {<ps_num>} {detail}`.

The output was modified to include power management details for modular switches in ExtremeXOS 11.3.

DC power output for the BlackDiamond 10808 switch and the BlackDiamond 8800 series switch was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all available platforms.

show power budget

```
show power budget
```

Description

Displays the power status and the amount of available and required power on a modular switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the amount of power available on the switch.

This status information may be useful if the `show slot` command displays a state of Powered OFF for any I/O module, for monitoring power, or for power planning purposes.

The first table of the `show power budget` command displays:

- Slot number of the power supply.
- Current state of the power supply. Options are:
 - Empty—There is no power supply installed.
 - Power Failed—The power supply has failed.
 - Power Off—The power supply is off.
 - Power On—The power supply is on.
- Watts and voltage amounts of the power supply.
- Redundant power information. Redundant power is the amount of power available if power to one PSU is lost. If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, the amount of redundant power shown is based on the worst-case assumption that power to a PSU with 220V AC input is lost.

The second table of the `show power budget` command displays:

- Slot number and name of the component installed in the slot. Options include:
 - I/O modules
 - MSM modules
 - Fan trays
- Current state of the module. Options include, among others:
 - Empty: There is no component installed.
 - Operational: The component is installed and operational.
 - Present: The component is installed but not operational.

- Down: The module is installed, but the administrator has taken the module offline.
- Power ON: There is sufficient system power to power up the module.
- Powered OFF: There is insufficient system power to keep the module up and running, or there is a mismatch between the module configured for the slot and the actual module installed in the slot.
- Booting: The module has completed downloading the software image and is now booting.
- Initializing: The module is initializing.
- Watts and voltage amounts of the modules.
- Power Surplus or Power Shortfall.
 - If the amount of available power meets or exceeds the required port, the excess is displayed as the Power Surplus.
 - If the available power is insufficient to meet the required power, the deficit is displayed as Power Shortfall.
- Redundant power information. If the amount of redundant power meets or exceeds the required power, the system has (N+1) power.
 - Yes—The system has redundant (N+1) power.
 - No—The system does not have redundant (N+1) power.

The information contained in this display is for planning purposes since the system operates without redundant power as long as a power surplus is shown. However, if power is lost to a single PSU when the system is not redundant, I/O modules are powered down. Please refer to the section [“Understanding Power Supply Management”](#) in [Chapter 4, “Managing the Switch,”](#) of the *ExtremeXOS Concepts Guide*.

Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different power information may be displayed.

Example

The following command displays the distribution of power and the available power on the switch:

```
show power budget
```

The following is sample output of this command from a BlackDiamond 12800 series switch:

```
PS   State                               48V
-----
 1   Powered On                           624.00
 2   Powered On                           624.00
 3   Empty
 4   Empty
 5   Empty
 6   Empty
-----
Power Available:                          1248.00
Redundant (N+1) Power Available:          648.00

Slots   Type           State           Watts
-----
Slot-1
Slot-2  GM-20T         Operational    149.00
Slot-5  GM-20T         Operational    149.00
```

Slot-6	Empty	
MSM-A	MSM-5	Operational 185.00
MSM-B	Empty	185.00
FanTray	Operational	45.00

Power Required:		713.00
Power Allocated:		713.00
Power Surplus:		535.00
Redundant Power Supply(s) Present?: NO		

History

This command was first available in ExtremeXOS 11.0.

Power over Ethernet (PoE) data (inline power) was added to the `show power budget` output in ExtremeXOS 11.1. PoE data is displayed when you install a G48P module in the BlackDiamond 8800 series switch.

Redundant (N+1) power information was added to the `show power budget` output in ExtremeXOS 11.1.

Platform Availability

This command is available only on modular switches.

show power controller

```
show power controller {<num>}
```

Description

Displays the current status of the installed power supply controllers.

Command Syntax

num	Specifies the slot number of the installed power supply controller.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supply controllers. Power controllers collect data about the installed power supplies and report the results to the MSM.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects the following power supply controller information:

- State—Indicates the current state of the power supply controller. Options are:
 - Empty: There is no power supply controller installed.
 - Operational: The power supply controller is installed and operational.
 - Present: The power supply controller is installed.
- PartInfo—Provides information about the power supply controller including the:
 - Slot number where the power supply controller is installed.
 - Serial number, a collection of numbers and letters, that make up the serial number of the power supply controller.
 - Part number, a collection of numbers and letters that make up the part number of the power supply controller.
- Revision—Displays the revision number of the power supply controller.
- FailureCode—Specifies the failure code of the power supply controller.
- Odometer—Specifies the date and how long the power supply controller has been operating.
- Temperature—Specifies, in celsius, the current temperature of the power supply controller.
- Status—Specifies the status of the power supply controller.

Example

The following command displays the status of the installed power supply controllers:

```
show power controller
```

The following is sample output from this command:

```
PSUCTRL-1 information:
  State:          Operational
  PartInfo:       PSUCTRL-1 04334-00021 450117-00-01
  Revision:       1.0
  FailureCode:    0
  Odometer:       337 days 7 hours   since Nov-30-2004
  Temperature:    32.14 deg C
  Status:         PSU CTRL Mode:   Master
```

```
PSUCTRL-2 information:
  State:          Empty
```

If you have two power supply controllers installed, the switch displays output about both of the power supply controllers:

```
PSUCTRL-1 information:
  State:          Operational
  PartInfo:       PSUCTRL-1 04334-00021 450117-00-01
  Revision:       1.0
  FailureCode:    0
  Odometer:       17 days 5 hours 30 minutes   since Oct-19-2004
  Temperature:    35.1 deg C
  Status:         PSU CTRL Mode:   Master
```

```
PSUCTRL-2 information:
  State:          Operational
  PartInfo:       PSUCTRL-2 04334-00068 450117-00-01
  Revision:       1.0
  FailureCode:    0
  Odometer:       4 days 13 hours   since Sep-21-2004
  Temperature:    33.56 deg C
  Status:         PSU CTRL Mode:   Backup
```

History

This command was first available in an ExtremeXOS 11.0.

Platform Availability

This command is available only on modular switches.

show session

```
show session [{detail} {<sessID>}] {history}
```

Description

Displays the currently active Telnet and console sessions communicating with the switch.

Syntax Description

detail	Specifies more detailed session information.
sessID	Specifies a session ID number.
history	Displays a list of all sessions.

Default

N/A.

Usage Guidelines

The `show session` command displays the username and IP address of the incoming Telnet session, whether a console session is currently active, and the login time. Each session is numbered.

Beginning with ExtremeXOS 11.2, the switch accepts IPv6 connections. If the incoming session is from an IPv6 address, the `show session` output indicates IPv6.

You can specify the following options to alter the session output:

- `detail`—The output for all current sessions is displayed in a list format.
- `sessID`—The output for the specified session is displayed in a list format.
- `history`—Displays a list of current and previous sessions, including the user, type of session, location, and start and end time of the session.

The `show session` command fields are defined in [Table 11](#).

Table 11: Show command field definitions

Field	Definition
#	Indicates session number.
Login Time	Indicates login time of session.
User	Indicates the user logged in for each session.
Type	Indicates the type of session, for example: console, telnet, http, https.
Auth	Indicates how the user is logged in.
CLI Auth	Indicates the type of authentication (RADIUS and TACAS) if enabled.
Location	Indicates the location (IP address) from which the user logged in. The output also indicates if the location is an IPv6 address.

Example

The following command displays the active sessions on the switch:

```
show session
```

The following is sample output from this command:

#	Login Time	User	Type	Auth	Auth Location	CLI
1	Thu Apr 28 20:16:56 2005	admin	console	local	dis	serial
*2	Thu Apr 28 23:36:20 2005	admin	ssh2	local	dis	
3001::20d:88ff:fec5:ad40						
3	Fri Apr 29 11:14:27 2005	admin	telnet	local	dis	10.255.44.55

The following command displays a list of current and previous sessions on the switch:

```
show session history
```

The following is sample output from this command:

```
Session History:
admin          console    serial          Mon Jun 21 09:19:
00 2004 Mon Jun 21 10:00:16 2004
admin          console    serial          Tue Jun 22 07:28:
11 2004 Tue Jun 22 11:46:48 2004
admin          console    serial          Wed Jun 23 10:05:
44 2004 Wed Jun 23 14:11:47 2004
admin          console    serial          Thu Jun 24 07:07:
25 2004 Thu Jun 24 07:08:55 2004
admin          console    serial          Thu Jun 24 13:30:
07 2004 Active
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show snmpv3 access

```
show snmpv3 access {[[hex <hex_group_name>] | <group_name>]}
```

Description

Displays SNMPv3 access rights.

Syntax Description

hex	Specifies that the value to follow is to be supplied as a colon separated string of hex octets.
group_name	Specifies the name of the group to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 access` command displays the access rights of a group. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmAccessTable entries.

Example

The following command displays all the access details:

```
show snmpv3 access
```

The following is sample output from this command:

```
Group Name       : admin
Context Prefix   :
Security Model    : USM
Security Level    : Authentication Privacy
Context Match     : Exact
Read View        : defaultAdminView
Write View        : defaultAdminView
Notify View       : defaultNotifyView
Storage Type      : Permanent
Row Status        : Active

Group Name       : initial
Context Prefix    :
Security Model    : USM
Security Level    : No-Authentication No-Privacy
Context Match     : Exact
Read View         : defaultUserView
```

```

Write View      :
Notify View     : defaultNotifyView
Storage Type    : Permanent
Row Status      : Active

Group Name      : v1v2c_ro
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View       :
Notify View      : defaultNotifyView
Storage Type     : Permanent
Row Status      : Active

Group Name      : v1v2c_rw
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View       : defaultUserView
Notify View      : defaultNotifyView
Storage Type     : Permanent
Row Status      : Active

Group Name      : v1v2cNotifygroup
Context Prefix  :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       :
Write View       :
Notify View      : defaultNotifyView
Storage Type     : Permanent
Row Status      : Active

Group Name      : v1v2cNotifyGroup
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       :
Write View       :
Notify View      : defaultNotifyView
Storage Type     : Permanent
Row Status      : Active

```

Total num. of entries in vacmAccessTable : 6

The following command displays the access rights for the group *group1*:

```
show snmpv3 access group1
```


History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 context

```
show snmpv3 context
```

Description

Displays information about the SNMPv3 contexts on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

This command displays the entries in the View-based Access Control Model (VACM) context table (VACMContextTable).

Example

The following command displays information about the SNMPv3 contexts on the switch:

```
show snmpv3 context
```

The following is sample output from this command:

```
VACM Context Name :  
Note : This Version Supports one global context ("")
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 counters

```
show snmpv3 counters
```

Description

Displays SNMPv3 counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show snmpv3 counters` command displays the following SNMPv3 counters:

- snmpUnknownSecurityModels
- snmpInvalidMessages
- snmpUnknownPDUHandlers
- usmStatsUnsupportedSecLevels
- usmStatsNotInTimeWindows
- usmStatsUnknownUserNames
- usmStatsUnknownEngineIDs
- usmStatsWrongDigests
- usmStatsDecryptionErrors

Issuing the command `clear counters` resets all counters to zero.

Example

The following command displays all the SNMPv3 counters.

```
show snmpv3 counters
```

The following is sample output from this command:

```
snmpUnknownSecurityModels      : 0
snmpInvalidMessages            : 0
snmpUnknownPDUHandlers         : 0
usmStatsUnsupportedSecLevels    : 0
usmStatsNotInTimeWindows       : 0
usmStatsUnknownUserNames       : 0
usmStatsUnknownEngineIDs       : 0
usmStatsWrongDigests           : 0
usmStatsDecryptionErrors       : 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 engine-info

```
show snmpv3 engine-info
```

Description

Displays information about the SNMPv3 engine on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

The following show engine-info output is displayed:

- Engine-ID—Either the ID auto generated from MAC address of switch, or the ID manually configured.
- Engine Boots—Number of times the agent has been rebooted.
- Engine Time—Time since agent last rebooted, in centiseconds.
- Max. Message Size—Maximum SNMP Message size supported by the Engine (8192).

Example

The following command displays information about the SNMPv3 engine on the switch:

```
show snmpv3 engine-info
```

The following is sample output from this command:

```
SNMP Engine-ID       : 80:0:7:7c:3:0:30:48:41:ed:97 'H'
SNMP Engine Boots    : 1
SNMP Engine Time     : 866896
SNMP Max. Message Size : 8192
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 filter

```
show snmpv3 filter {[[hex <hex_profile_name>] | <profile_name>] {{subtree}
<object_identifier>}
```

Description

Displays the filters that belong a filter profile.

Syntax Description

hex_profile_name	Specifies the filter profile to display. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile to display in ASCII format.
object_identifier	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to display entries from the snmpNotifyFilterTable. If you specify a profile name and subtree, you will display only the entries with that profile name and subtree. If you specify only the profile name, you will display all entries for that profile name. If you do not specify a profile name, then all the entries are displayed.

Example

The following command displays the part of filter profile *prof1* that includes the MIB subtree *1.3.6.1.4.1*:

```
show snmpv3 filter prof1 subtree 1.3.6.1.4.1
```

The following is sample output from this command:

```
Profile Name      : prof1
Subtree           : 1.3.6.1.4.1
Mask              :
Type              : Included
Storage Type      : NonVolatile
Row Status        : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 filter-profile

```
show snmpv3 filter-profile {[[hex <hex_profile_name>] | <profile_name>]}
{param [[hex <hex_param_name>] | <param_name>]}
```

Description

Displays the association between parameter names and filter profiles.

Syntax Description

hex_profile_name	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name in ASCII format.
hex_param_name	Specifies the parameter name. The values is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command displays the entry with filter profile *prof1* with the parameter name *P1*:

```
show snmpv3 filter-profile prof1 param P1
```

The following is sample output of this command:

```
Filter Profile Params Name : p1
Name                       : prof1
Storage Type               : NonVolatile
Row Status                 : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 group

```
show snmpv3 group {[[hex <hex_group_name>] | <group_name>] {user [[hex
<hex_user_name>] | <user_name>]]}}
```

Description

Displays the user name (security name) and security model association with a group name.

Syntax Description

hex_group_name	Specifies the group name to display. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to display. The value is to be supplied in ASCII format.
hex_user_name	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 group` command displays the details of a group with the given group name. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmSecurityToGroupTable.

Example

The following command displays information about all groups for every security model and user name:

```
show snmpv3 group
```

The following is sample output from this command:

```
Group Name       : v1v2c_ro
Security Name    : v1v2c_ro
Security Model    : snmpv1
Storage Type     : Permanent
Row Status       : Active

Group Name       : v1v2c_rw
Security Name    : v1v2c_rw
Security Model    : snmpv1
Storage Type     : Permanent
Row Status       : Active
```

```

Group Name       : v1v2c_ro
Security Name    : v1v2c_ro
Security Model    : snmpv2c
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : v1v2c_rw
Security Name    : v1v2c_rw
Security Model    : snmpv2c
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : admin
Security Name    : admin
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : initial
Security Name    : initial
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : initial
Security Name    : initialmd5
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : initial
Security Name    : initialsha
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : initial
Security Name    : initialmd5Priv
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```

Group Name       : initial
Security Name    : initialshaPriv
Security Model    : USM
Storage Type     : Permanent
Row Status       : Active

```

```
Total num. of entries in vacmSecurityToGroupTable : 10
```

The following command shows information about the group *testgroup* and user name *testuser*:

```
show snmpv3 group testgroup user testuser
```

The following is sample output from this command:

```
Group Name       : testgroup
Security Name    : testuser
Security Model   : USM
Storage Type     : NonVolatile
Row Status       : Active
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_group_name` and `hex_user_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 mib-view

```
show snmpv3 mib-view {[[hex <hex_view_name>] | <view_name>] {subtree
<object_identifier>}}
```

Description

Displays a MIB view.

Syntax Description

hex_view_name	Specifies the name of the MIB view to display. The value is to be supplied as a colon separated string of hex octets.
view_name	Specifies the name of the MIB view to display. The value is to be supplied in ASCII format.
object_identifier	Specifies the object identifier of the view to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 mib-view` command displays a MIB view. If you do not specify a view name, the command will display details for all the MIB views. If a subtree is not specified, then all subtrees belonging to the view name will be displayed.

This command displays the SNMPv3 vacmViewTreeFamilyTable.

Example

The following command displays all the view details:

```
show snmpv3 mib-view
```

The following is sample output from this command:

```
View Name       : defaultUserView
MIB Subtree     : 1
View Type       : Included
Storage Type    : Permanent
Row Status      : Active
```

```
View Name       : defaultUserView
MIB Subtree     : 1.3.6.1.6.3.16
View Type       : Excluded
Storage Type    : Permanent
Row Status      : Active
```

```

View Name      : defaultUIView
MIB Subtree    : 1.3.6.1.6.3.18
View Type      : Excluded
Storage Type   : Permanent
Row Status     : Active

View Name      : defaultUIView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.4
View Type      : Excluded
Storage Type   : Permanent
Row Status     : Active

View Name      : defaultUIView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.6
View Type      : Excluded
Storage Type   : Permanent
Row Status     : Active

View Name      : defaultUIView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.9
View Type      : Excluded
Storage Type   : Permanent
Row Status     : Active

View Name      : defaultAdminView
MIB Subtree    : 1
View Type      : Included
Storage Type   : Permanent
Row Status     : Active

View Name      : defaultNotifyView
MIB Subtree    : 1
View Type      : Included
Storage Type   : Permanent
Row Status     : Active

```

Total num. of entries in vacmViewTreeFamilyTable : 8

The following command displays a view with the view name *Roview* and subtree 1.3.6.1.2.1.1:

```
show snmpv3 mib-view Roview subtree 1.3.6.1.2.1.1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 notify

```
show snmpv3 notify {[[hex <hex_notify_name>] | <notify_name>]}
```

Description

Displays the notifications that are set. This command displays the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the parameter name associated with the target. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpNotifyTable. This table lists the notify tags that the agent will use to send notifications (traps).

If no notify name is specified, all the entries are displayed.

Example

The following command displays the notify table entry for *N1*:

```
show snmpv3 notify N1
```

The following is sample output from this command:

```
Notify Name      : N1
Tag              : type1
Type             : Trap
Storage Type     : NonVolatile
Row Status       : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 target-addr

```
show snmpv3 target-addr {[[hex <hex_addr_name>] | <addr_name>]}
```

Description

Displays information about SNMPv3 target addresses.

Syntax Description

hex_addr_name	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 snmpTargetAddressTable. If no target address is specified, the entries for all the target addresses will be displayed.

To view the source IP address, use the [show management](#) command.

Example

The following command displays the entry for the target address named *A1*:

```
show snmpv3 target-addr A1
```

The following is sample output from this command:

```
Target Addr Name      : A1
TDomain               : 1.3.6.1.6.1.1
TAddress              : 10.201.31.234, 162
TMask                 :
Timeout               : 1500
Retry Count           : 0
Tag List              : defaultNotify
Params                : v1v2cNotifyParam1
Storage Type          : NonVolatile
Row Status             : Active
Storage Type          : NonVolatile
Row Status             : Active
```


History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 11.0 to display a list of tags if more than one was configured and to display the timeout value for the entry in the snmpTargetAddrTable. This command was also modified to support the `hex_addr_name` parameter.

Platform Availability

This command is available on all platforms.

show snmpv3 extreme-target-addr-ext

```
show snmpv3 extreme-target-addr-ext [[hex <hex_addr_name>] | <addr_name>]
```

Description

Displays information about SNMPv3 target addresses enhanced or standard mode.

Syntax Description

hex_addr_name	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 extremeTargetAddressExtTable.

Example

The following command displays the entry for the target address named *A1*:

```
show snmpv3 extreme-target-addr-ext A1
```

The following is sample output from this command:

```
Target Addr Name      : A1
Mode                  : Enhanced
IgnoreMPModel         : No
UseEventComm          : Yes
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_addr_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 target-params

```
show snmpv3 target-params {[[hex <hex_target_params>] | <target_params>]}
```

Description

Displays the information about the options associated with the parameter name.

Syntax Description

hex_target_params	Specifies the parameter to display. The value is to be supplied as a colon separated string of hex octets.
target_params	Specifies the parameter name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

If no parameter name is specified, all the entries are displayed.

Example

The following command displays the target parameter entry named *P1*:

```
show snmpv3 target-params P1
```

The following is sample output from this command:

```
Target Params Name      : p1
MP Model                : snmpv2c
Security Model          : snmpv2c
User Name               : testuser
Security Level          : No-Authentication No-Privacy
Storage Type            : NonVolatile
Row Status              : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_target_params parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 user

```
show snmpv3 user {[[hex <hex_user_name>] | <user_name>]}
```

Description

Displays detailed information about the user.

Syntax Description

hex_user_name	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 user` command displays the details of a user. If you do not specify a user name, the command will display details for all the users. The authentication and privacy passwords and keys will not be displayed.

The user entries in SNMPv3 are stored in the USMUserTable, so the entries are indexed by EngineID and user name.

Example

The following command lists all user entries:

```
show snmpv3 user
```

The following is sample output from this command:

```
Engine-ID       : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name       : admin
Security Name   : admin
Authentication   : HMAC-MD5
Privacy         : DES
Storage Type    : Permanent
Row Status      : Active

Engine-ID       : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name       : initial
Security Name   : initial
Authentication   : No-Authentication
Privacy         : No-Privacy
```

```

Storage Type      : Permanent
Row Status        : Active

Engine-ID         : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name         : initialmd5
Security Name     : initialmd5
Authentication    : HMAC-MD5
Privacy           : No-Privacy
Storage Type      : Permanent
Row Status        : Active

```

```

Engine-ID         : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name         : initialsha
Security Name     : initialsha
Authentication    : HMAC-SHA
Privacy           : No-Privacy
Storage Type      : Permanent
Row Status        : Active

```

```

Engine-ID         : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name         : initialmd5Priv
Security Name     : initialmd5Priv
Authentication    : HMAC-MD5
Privacy           : DES
Storage Type      : Permanent
Row Status        : Active

```

```

Engine-ID         : 80:00:07:7c:03:00:04:96:1f:a2:60 'H'
User Name         : initialshaPriv
Security Name     : initialshaPriv
Authentication    : HMAC-SHA
Privacy           : DES
Storage Type      : Permanent
Row Status        : Active

```

Total num. of entries in usmUserTable : 6

The following command lists details for the specified user, *testuser*:

```
show snmpv3 user testuser
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show sntp-client

```
show sntp-client
```

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays configuration and statistics information of SNTP client.

Example

The following command displays the SNTP configuration:

```
show sntp-client
```

The following is sample output from this command:

```
SNTP client is enabled
SNTP time is valid
Primary server: 172.17.1.104
Secondary server: 172.17.1.104
Query interval: 64
Last valid SNTP update: From server 172.17.1.104, on Wed Oct 30 22:46:03 2003
SNTPC Statistics:
Packets transmitted:
  to primary server:          1
  to secondary server:        0
Packets received with valid time:
  from Primary server:        1
  from Secondary server:      0
  from Broadcast server:      0
Packets received without valid time:
  from Primary server:        0
  from Secondary server:      0
  from Broadcast server:      0
Replies not received to requests:
  from Primary server:        0
  from Secondary server:      0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

telnet

```
telnet {vr <vr_name>} [<host_name> | <remote_ip>] [<port>]
```

Description

Allows you to Telnet from the current command-line interface session to another host.

Syntax Description

vr	Specifies use of a virtual router. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created virtual routers.
vr_name	Specifies the name of the virtual router.
host_name	Specifies the name of the host.
remote_ip	Specifies the IP address of the host.
port	Specifies a TCP port number. The default is port 23.

Default

- Telnet—enabled
- Virtual router—Uses all virtual routers on the switch for outgoing Telnet requests
- Port—23

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created virtual routers.

Only VT100 emulation is supported.

Before you can start an outgoing Telnet session, you need to configure the switch IP parameters. To open a Telnet connection, you must specify the host IP address or the host name of the device you want to connect to. Check the user manual supplied with the Telnet facility if you are unsure of how to do this. Although the switch accepts IPv6 connections, you can only Telnet from the switch to another device with an IPv4 address.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Virtual Router Requirements. The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are VR-Mgmt, VR-Control, and VR-Default; however, you can only Telnet on VR-Mgmt and VR-Default. On the BlackDiamond 10808 switch, you can also create and configure your own virtual routers. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively, and Telnet used VR-0 by default. For more information about virtual routers, see [“Virtual Routers”](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command starts a Telnet client communication to the host at IP address 123.45.67.8:

```
telnet 123.45.67.8
```

The following command starts a Telnet client communication with a host named *sales*:

```
telnet sales
```

History

This command was first available in ExtremeXOS 10.1.

Support for the following virtual routers was added in ExtremeXOS 11.0: VR-Mgmt and VR-Default.

Platform Availability

This command is available on all platforms.

telnet msm

```
telnet msm [a | b]
```

Description

Allows you to Telnet to either the primary or the backup MSM regardless of which console port you are connected to.

Syntax Description

a	Specifies the MSM installed in slot A.
b	Specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command to access either the primary or the backup MSM regardless of which console port you are connected to. For example, if MSM A is the primary MSM and you are connected to MSM A via its console port, you can access the backup MSM installed in slot B by issuing the `telnet msm b` command.

Example

The following example makes the following assumptions:

- The MSM installed in slot A is the primary
- The MSM installed in slot B is the backup
- You have a console connection to MSM B

The following command accesses the primary MSM installed in slot A from the backup MSM installed in slot B:

```
My8800.6 # telnet msm b
```

```
Entering character mode
Escape character is '^['.
```

```
telnet session telnet0 on /dev/ptyb0
```

```
login: admin
password:
```

```
ExtremeXOS
Copyright (C) 2000-2007 Extreme Networks. All rights reserved.
```

Protected by US Patent Nos: 6,678,248; 6,104,700; 6,766,482; 6,618,388; 6,034,957;
 6,859,438; 6,912,592; 6,954,436; 6,977,891; 6,980,550; 6,981,174; 7,003,705;
 7,012,082; 7,046,665; 7,126,923; 7,142,509; 7,149,217; 7,152,124; 7,154,861.
 =====

You are connected to a Backup node. Only a limited command set is supported.
 You may use "telnet msm A" to connect to the Master node to access
 the full set of commands.

Press the <tab> or '?' key at any time for completions.
 Remember to save your configuration changes.

My8800.1 >

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on modular switches.

telnet slot

```
telnet slot <slot-number>
```

Description

Allows you to Telnet from any source node to any other target node when both nodes are in the same active topology.

Syntax Description

slot-number	Specifies the number of a slot that is currently occupied by the intended target node.

Default

N/A.

Usage Guidelines

When the target node accepts the related TCP connection, it prompts the user for a user ID and password. If the failsafe account is used, user ID and password authentication takes place on the specified node. Otherwise, authentication takes place on the master node, regardless of the source and target nodes used.

Telnet must be enabled on the SummitStack for this command to function.

Example

The following command accesses the node in slot 2:

```
Slot-1 Stack.6 # telnet slot 2
```

```
Entering character mode
Escape character is '^]'.
```

```
telnet session telnet0 on /dev/ptyb0
```

```
login: admin
password:
```

```
ExtremeXOS
Copyright (C) 2000-2007 Extreme Networks. All rights reserved.
Protected by US Patent Nos: 6,678,248; 6,104,700; 6,766,482; 6,618,388; 6,034,957;
6,859,438; 6,912,592; 6,954,436; 6,977,891; 6,980,550; 6,981,174; 7,003,705;
7,012,082; 7,046,665; 7,126,923; 7,142,509; 7,149,217; 7,152,124; 7,154,861.
=====
```

```
You are connected to a Backup node. Only a limited command set is supported.
You may use "telnet slot <slot_number>" to connect to the Master node to access
```

the full set of commands.

Press the <tab> or '?' key at any time for completions.
Remember to save your configuration changes.

```
Slot-2 Stack.1 >
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack.

tftp

```
tftp [<host-name> | <ip-address>] {-v <vr_name>} [-g | -p] [{-l [internal-memory <local-file-internal> | memorycard <local-file-memcard> | <local-file>}] {-r <remote-file>}] | {-r <remote-file>} {-l [internal-memory <local-file-internal> | memorycard <local-file-memcard> | <local-file>]]}]
```

Description

Allows you to TFTP from the current command line interface session to a TFTP server.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
-g	Gets the specified file from the TFTP server and copies it to the local host.
-p	Puts the specified file from the local host and copies it to the TFTP server.
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card. NOTE: This parameter is available only on modular switches.
local-file-memcard	Specifies the name of the file on the external compact flash card. NOTE: This parameter is available only on modular switches.
local-file	Specifies the name of the file (configuration file, policy file) on the local host.
remote-file	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

NetASCII and mail file type formats are not supported.

TFTP Server Requirements. Extreme Networks recommends using a TFTP server that supports blocksize negotiation (as described in RFC 2348, *TFTP Blocksize Option*), to enable faster file downloads and larger file downloads. If the TFTP server does not support blocksize negotiation, the file size is limited to

32 MB. Older TFTP servers that do not support blocksize negotiation have additional implementation limits that may decrease the maximum file size to only 16 MB, which may be too small to install ExtremeXOS images.

If your TFTP server does not support blocksize negotiation, the switch displays a message similar to the following when you attempt a get (-g) or put (-p) operation:

```
Note: The blocksize option is not supported by the remote TFTP server.
      Without this option, the maximum file transfer size is limited to 32MB.
      Some older TFTP servers may be limited to 16MB file.
```

Using TFTP. Use TFTP to download a previously saved configuration file or policy file from the TFTP server to the switch. When you download a file, this command does not automatically apply it to the switch. You must specify that the downloaded file be applied to the switch. For example, if you download a configuration file, issue the `use configuration` command to apply the saved configuration on the next reboot. You must use the `reboot` command to activate the new configuration. If you download a policy file, use the `refresh policy` command to reprocess the text file and update the policy database.

You also use TFTP to upload a saved configuration file or policy file from the switch to the TFTP server.

If your download from the TFTP server to the switch is successful, the switch displays a message similar to the following:

```
Downloading megtest2.cfg to switch... done!
```

If your upload from the switch to the TFTP server is successful, the switch displays a message similar to the following:

```
Uploading megtest1.cfg to TFTPHost ... done!
```

Up to eight active TFTP sessions can run on the switch concurrently.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Virtual Router Requirements. The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are VR-Mgmt, VR-Control, and VR-Default; however, you can only TFTP on VR-Mgmt and VR-Default. On the BlackDiamond 10808 switch, you can also create and configure your own virtual routers. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively. For more information about virtual routers, see [“Virtual Routers”](#) in the *ExtremeXOS Concepts Guide*.

Internal Memory and Core Dump Files. Core dump files have a .gz file extension. The filename format is: `core.<process-name.pid>.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you have a modular switch and save core dump files to the external memory card, the filename also includes the affected MSM: MSM-A or MSM-B.

If you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to transfer those files from the internal memory card to a TFTP server. If you have a modular switch, you can transfer core dump information to and from the external compact flash card.

If the switch has not saved any debug files, you cannot transfer other files to or from the internal memory. For example if you attempt to transfer a configuration file from the switch to the internal memory, the switch displays a message similar to the following:

```
Error: tftp transfer to internal-memory not allowed.
```

For information about configuring and sending core dump information to the internal memory card, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see [Chapter B, “Troubleshooting,”](#) in the *ExtremeXOS Concepts Guide*.

Modular Switches Only. If you specify the `memorycard` option, you can copy and transfer files to and from the external memory card using TFTP.

Other Useful Commands. On the Summit family switches, SummitStack, and the BlackDiamond 10808 switch, use the [download bootrom](#) command to upgrade the BootROM. This command utilizes TFTP to transfer the BootROM image file from your TFTP server to the switch. Only upgrade the BootROM when asked to do so by an Extreme Networks technical representative. For more information about this command, see [“download bootrom” on page 2365](#).

To upgrade the image, use the [download image](#) command. This command utilizes TFTP to transfer the software image file from your TFTP server to the switch. For more information about this command, see [“download image” on page 2368](#).

Example

The following command downloads the configuration file named *XOS1.cfg* from the TFTP server with an IP address of 10.123.45.67:

```
tftp 10.123.45.67 -g -r XOS1.cfg
```

The following command retrieves and transfers files from the external memory card installed in a modular switch:

```
tftp 10.1.2.3. -g -l memorycard test.pol -r august23.pol
```

History

This command was first available in ExtremeXOS 10.1.

The `memorycard` option was added in ExtremeXOS 11.1.

The `internal-memory` option was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

tftp get

```
tftp get [<host-name> | <ip-address>] {-vr <vr_name>} [[[internal-memory
<local-file-internal> | memorycard <local-file-memcard> | <local_file>]
{<remote_file>} | {<remote_file>} {[internal-memory <local-file-internal> |
memorycard <local-file-memcard> | <local_file>}]]] {force-overwrite}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from a TFTP server and copy it to a local host, including the switch, internal memory card, or external compact flash card.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card. NOTE: This parameter is available only on modular switches.
local-file-memcard	Specifies the name of the file on the external compact flash card. NOTE: This parameter is available only on modular switches.
local_file	Specifies the name of the file (configuration file, policy file) on the local host.
remote_file	Specifies the name of the file on the remote host.
force-overwrite	Specifies the switch to automatically overwrite an existing file.

Default

If you do not specify a virtual router, VR-Mgmt is used; if you transfer a file with a name that already exists on the system, the switch prompts you to overwrite the existing file.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

NetASCII and mail file type formats are not supported.

By default, the switch prompts you to overwrite an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch displays a message similar to the following:

```
test.cfg already exists, do you want to overwrite it? (y/n)
```

Enter `y` to download the file and overwrite the existing file. Enter `n` to cancel this action.

If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

If you cancel this action, the switch displays a message similar to the following:

```
Tftp download aborted.
```

If you specify the `force-overwrite` parameter, the switch automatically overwrites an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch automatically overrides the existing file. If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original `tftp` command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements
- How to use TFTP
- Host name and remote IP address character restrictions
- Local and remote filename character restrictions
- Virtual router requirements
- Internal memory and core dump files
- Topics related to only modular switches
- Other useful commands

See the `tftp` command [on page 259](#).

Example

The following command retrieves and transfers the file `test.pol` from a TFTP server with an IP address of `10.1.2.3` and renames the file `august23.pol` when transferred to the external memory card installed in a modular switch:

```
tftp get 10.1.2.3 test.pol memory-card august23.pol
```

The following command retrieves the configuration file named `meg-upload.cfg` from a TFTP server with an IP address of `10.10.10.10`:

```
tftp get 10.10.10.10 meg_upload.cfg
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

tftp put

```
tftp put [<host-name> | <ip-address>] {-vr <vr_name>} [[[internal-memory
<local-file-internal> | memorycard <local-file-memcard> | <local_file>}
{<remote_file>} | {<remote_file>} {[internal-memory <local-file-internal> |
memorycard <local-file-memcard> | <local_file>}]]]
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from the local host, including the switch, internal memory card, or external compact flash card and put it on a TFTP server.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card. NOTE: This parameter is available only on modular switches.
local-file-memcard	Specifies the name of the file on the external compact flash card. NOTE: This parameter is available only on modular switches.
local_file	Specifies the name of the file (configuration file, policy file) on the local host.
remote_file	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

NetASCII and mail file type formats are not supported.

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original TFTP command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements
- How to use TFTP
- Host name and remote IP address character restrictions
- Local and remote filename character restrictions
- Virtual router requirements
- Internal memory and core dump files
- Topics related to only modular switches
- Other useful commands

See the [tftp](#) command [on page 259](#).

Example

The following command transfers a saved, not currently used configuration file named *XOS1.cfg* from the switch to the TFTP server:

```
tftp put 10.123.45.67 XOS1.cfg
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

4

Commands for Managing the ExtremeXOS Software

This chapter describes commands for:

- Working with the configuration and policy files used by the switch
- Starting, stopping, and displaying information about processes on the switch
- Viewing system memory resources
- Monitoring CPU utilization

**NOTE**

For information about downloading and upgrading a new software image, saving configuration changes, and upgrading the BootROM, see [Appendix A, “Configuration and Image Commands.”](#)

Like any advanced operating system, ExtremeXOS gives you the tools to manage your switch and create your network configurations. With the introduction of ExtremeXOS, the following enhancements and functionality have been added to the switch operating system:

- File system administration—You can move, copy, and delete files from the switch. The file system structure allows you to keep, save, rename, and maintain multiple copies of configuration files on the switch. In addition, you can manage other entities of the switch such as policies and access control lists (ACLs).
- Configuration file management—You can oversee and manage multiple configuration files on your switch. In addition, you can upload, download, modify, and name configuration files used by the switch.
- Process control—You can stop and start processes, restart failed processes, and update the software for a specific process or set of processes.
- Memory protection—With memory protection, ExtremeXOS protects each process from every other process in the system. If one process experiences a memory fault, that process cannot affect the memory space of another process.
- CPU monitoring—You can monitor CPU utilization for Management Switch Fabric Modules (MSMs) and the individual processes running on the switch. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes.

**NOTE**

Filenames are case-sensitive. For information on filename restrictions, please refer to the specific command in the ExtremeXOS Command Reference Guide.

clear cpu-monitoring

```
clear cpu-monitoring {process <name>} {slot <slotid>}
```

Description

Clears, resets the CPU utilization history and statistics stored in the switch.

Syntax Description

name	Specifies the name of the process.
slotid	Specifies the slot number of the MSM module: • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

If you do not specify any keywords, this command clears the CPU utilization history for the entire switch, including processes, and resets the statistics to zero (0). On modular switches, this command also clears the CPU utilization history of the installed MSMs.

If you specify `process`, the switch clears and resets the CPU utilization history for the specified process.

Modular Switches Only. If you specify `slot`, the switch clears and resets the CPU utilization history for the specified MSM.

Example

The following command resets the CPU history and resets the statistics to 0 for the TFTP process running on the MSM installed in slot A of a modular switch:

```
clear cpu-monitoring process tftpd slot A
```

The following command resets the CPU history and resets statistics to 0 for the TFTP process running on a Summit family switch:

```
clear cpu-monitoring process tftpd
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

cp

```
cp [internal-memory <old-name-internal> internal-memory <new-name-internal>
| internal-memory <old-name-internal> memorycard <new-name-memorycard> |
memorycard <old-name-memorycard> memorycard <new-name-memorycard> |
memorycard <old-name-memorycard> <new-name> | <old-name> memorycard <new-
name-memorycard> | <old-name> <new-name>]
```

Description

Copies an existing configuration, policy, or if configured, core dump file stored in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
old-name-internal	Specifies the name of the core dump file located on the internal memory card that you want to copy.
new-name-internal	Specifies the name of the newly copied core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash memory card. NOTE: This parameter is available only on modular switches.
old-name-memorycard	Specifies the name of the file located on the external compact flash memory card that you want to copy. Depending on your switch configuration, you can have configuration, policy, or core dump files stored in this card. NOTE: This parameter is available only on modular switches.
new-name-memorycard	Specifies the name of the newly copied file located on the external compact flash memory card. NOTE: This parameter is available only on modular switches.
old-name	Specifies the name of the configuration or policy file that you want to copy.
new-name	Specifies the name of the newly copied configuration or policy file.

Default

N/A.

Usage Guidelines

Use this command to make a copy of an existing file before you alter or edit the file. By making a copy, you can easily go back to the original file if needed.

If you copy a configuration or policy file, remember the following:

- XML-formatted configuration files have the .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have the .xsf file extension. See [Appendix A, “Software Upgrade and Boot Options,”](#) in the *ExtremeXOS Concepts Guide* for more information.
- Policy files have the .pol file extension.
- Core dump files have the .gz file extension. See [“Internal Memory and Core Dump Files”](#) for more information.

When you copy a configuration or policy file from the system, make sure you specify the appropriate file extension. For example, if you want to copy a policy file, specify the filename and *.pol*.

When you copy a file on the switch, the switch displays a message similar to the following:

```
Copy config test.cfg to config test1.cfg on switch? (y/n)
```

Enter *y* to copy the file. Enter *n* to cancel this process and not copy the file.

If you enter *y*, the switch copies the file with the new name and keeps a backup of the original file with the original name. After the switch copies the file, use the *ls* command to display a complete list of files. In this example, the switch displays the original file named *test.cfg* and the copied file named *test_rev2.cfg*.

The following is sample output from the *ls* command:

```
...
-rw-r--r--    1 root    root      100980 Sep 23 09:16 test.cfg
-rw-r--r--    1 root    root      100980 Oct 13 08:47 test_rev2.cfg
...
```

If you enter *n*, the switch displays a message similar to the following:

```
Copy cancelled.
```

Case-sensitive Filenames. Filenames are case-sensitive. In this example, you have a configuration file named *Test.cfg*. If you attempt to copy the file with the incorrect case, for example *test.cfg*, the switch displays a message similar to the following:

```
Error: cp: /config/test.cfg: No such file or directory
```

Since the switch is unable to locate *test.cfg*, the file is not copied.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Internal Memory and Core Dump Files. Core dump files have a *.gz* file extension. The filename format is: *core.<process-name.pid>.gz* where *process-name* indicates the name of the process that failed and *pid* is the numerical identifier of that process. If you have a modular switch and save core dump files to the external memory card, the filename also includes the affected MSM: MSM-A or MSM-B.

By making a copy of a core dump file, you can easily compare new debug information with the old file if needed.

If you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the *internal-memory* option and associated internal-memory name options to copy an

existing core dump file. If you have a modular switch with an external compact flash memory card installed, you can copy the core dump file to that card.

For information about configuring and sending core dump information to the internal memory card, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see [Appendix B, “Troubleshooting,”](#) in the *ExtremeXOS Concepts Guide*.

Modular Switches Only. This command also replicates the action from the primary MSM to the backup MSM. For example, if you copy a file on the primary MSM, the same file is copied to the backup MSM.

For the `memorycard` option, the source and/or destination is the memorycard. You must mount the memory card for this operation to succeed. The `cp` command copies a file from the switch to the external memory card or a file already on the card. If you copy a file from the switch to the external memory card, and the new filename is identical to the source file, you do not need to re-enter the filename.

If you send core dump information to the external memory card, specify the `memorycard` option and associated `memorycard` name options to copy an existing core dump file. For more information about core dump files, see [“Internal Memory and Core Dump Files”](#) on page 273.

Example

The following command makes a copy of a configuration file named `test.cfg` and gives the copied file a new name of `test_rev2.cfg`:

```
cp test.cfg test_rev2.cfg
```

The following command makes a copy of a configuration file named `primary.cfg` from a modular switch to the external memory card with the same name, `primary.cfg`:

```
cp primary.cfg memorycard
```

The above command performs the same action as entering the following command:

```
cp primary.cfg memorycard primary.cfg
```

History

This command was first available in ExtremeXOS 11.0.

The `memorycard` option was added in ExtremeXOS 11.1.

The `internal-memory` option was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable cpu-monitoring

```
disable cpu-monitoring
```

Description

Disables CPU monitoring on the switch.

Command Syntax

This command has no arguments or variables.

Default

CPU monitoring is enabled and occurs every 20 seconds.

Usage Guidelines

Use this command to disable CPU monitoring on the switch.

This command does not clear the monitoring interval. Therefore, if you altered the CPU monitoring interval, this command does not return the CPU monitoring interval to 20 seconds. To return to the default frequency level, use the `enable cpu-monitoring {interval <seconds>} {threshold <percent>}` and specify 20 for the interval.

Example

The following command disables CPU monitoring on the switch:

```
disable cpu-monitoring
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable xml-mode

```
disable xml-mode
```

Description

Disables XML configuration mode on the switch.

Command Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable the XML configuration mode on the switch. XML configuration mode is not supported for end users.

See the command:

```
enable xml-mode
```

Example

The following command disables XML configuration mode on the switch:

```
disable xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable cpu-monitoring

```
enable cpu-monitoring {interval <seconds>} {threshold <percent>}
```

Description

Enables CPU monitoring on the switch.

Command Syntax

seconds	Specifies the monitoring interval, in seconds. The default is 20 seconds, and the range is 5 to 60 seconds.
threshold	Specifies the CPU threshold value. CPU usage is measured in percentages. The default is 60%, and the range is 0% to 100%.

Default

CPU monitoring is enabled and occurs every 20 seconds. The default CPU threshold value is 60%.

Usage Guidelines

CPU monitoring allows you to monitor the CPU utilization and history for all of the processes running on the switch. By viewing this history on a regular basis, you can see trends emerging and identify processes with peak utilization. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes before they become a problem.

To specify the frequency of CPU monitoring, use the `interval` keyword. Extreme Networks recommends the default setting for most network environments. If you enter a number lower than 20 seconds, CPU utilization may increase.

CPU usage is measured in percentages. By default, the CPU threshold value is 60%. If CPU utilization of a process exceeds 60% of the regular operating basis, the switch logs an error message specifying the process name and the current CPU utilization for the process. In addition, if configured with the guidance of Extreme Networks Customer Support, the switch can save this information in a coredump or tracefile. To modify the CPU threshold level, use the `threshold` keyword. The range is 0% to 100%.

Example

The following command enables CPU monitoring every 30 seconds:

```
enable cpu-monitoring interval 30
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable xml-mode

```
enable xml-mode
```

Description

Enables XML configuration mode on the switch.

Command Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XML configuration mode on the switch, however XML configuration mode is not supported for end users, and Extreme Networks strongly cautions you not to enable this mode. Use this command only under the direction of Extreme Networks.

If you inadvertently issue this command, the switch prompt will be changed by adding the text `(xml)` to the front of the prompt. If you see this mode indicator, please disable XML configuration mode by using the following command:

```
disable xml-mode
```

Example

The following command enables XML configuration mode on the switch:

```
enable xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

ls

```
ls {internal-memory | memorycard}
```

Description

Lists all configuration, policy, and if configured, core dump files in the system.

Syntax Description

internal-memory	Lists the core dump (debug) files that are present and saved in the internal memory card.
memorycard	Lists all of the files on the removable external compact flash memory card. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

If you use issue this command without any options, the output displays all of the configuration and policy files stored on the switch.

If you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to display the core dump files stored on the internal memory card. For more information about core dump files, see [“Core Dump Files” on page 280](#).

If you specify the `memorycard` option on a modular switch, the output displays all of the files stored on the external compact flash memory card, including core dump files if so configured. For more information about core dump files, see [“Core Dump Files” on page 280](#).

Understanding the Output. Output from this command includes the following:

- The first column displays the file permission using the following ten place holders:
 - The first place holder displays - for a file.
 - The next three place holders display `r` for read access and `w` for write access permission for the file owner.
 - The following three place holders display `r` for read access permission for members of the file owner’s group.
 - The last three place holders display `r` for read access for every user that is not a member of the file owner’s group.
- The second column displays how many links the file has to other files or directories.
- The third column displays the file owner.
- The remaining columns display the file size, date and time the file was last modified, and the file name.

Core Dump Files. Core dump files have a .gz file extension. The filename format is: core.<process-name.pid>.gz where process-name indicates the name of the process that failed and pid is the numerical identifier of that process. If you have a modular switch and save core dump files to the external memory card, the filename also includes the affected MSM: MSM-A or MSM-B.

If the switch has not saved any debug files, no files are displayed. For information about configuring and sending core dump information to the internal memory card or the external memory card, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see [Appendix B, “Troubleshooting,”](#) in the *ExtremeXOS Concepts Guide*.

BlackDiamond 10808 Switch Only. In ExtremeXOS 10.1, this command was available only on the primary Management Switch Fabric Module (MSM). You were unable to display a list of configuration and policy files on the backup MSM.

Example

The following command displays a list of all current configuration and policy files in the system:

```
ls
```

The following is sample output from this command:

```
total 424
-rw-r--r--    1 root    root          50 Jul 30 14:19 hugh.pol
-rw-r--r--    1 root    root       94256 Jul 23 14:26 hughtest.cfg
-rw-r--r--    1 root    root     100980 Sep 23 09:16 megtest.cfg
-rw-r--r--    1 root    root         35 Jun 29 06:42 newpolicy.pol
-rw-r--r--    1 root    root     100980 Sep 23 09:17 primary.cfg
-rw-r--r--    1 root    root       94256 Jun 30 17:10 roytest.cfg
```

On a modular switch, the following command displays a list of all current configuration and policy files in the external memory card:

```
ls memorycard
```

The following is sample output from this command:

```
-rwxr-xr-x    1 root    0      15401865 Mar 30 00:03 bd10K-11.2.0.13.xos
-rwxr-xr-x    1 root    0          10 Mar 31 09:41 test-1.pol
-rwxr-xr-x    1 root    0          10 Apr  4 09:15 test.pol
-rwxr-xr-x    1 root    0          10 Mar 31 09:41 test_1.pol
-rwxr-xr-x    1 root    0     223599 Mar 31 10:02 v11_1_3.cfg
```

History

This command was first available in ExtremeXOS 10.1.

The `memorycard` option was added in ExtremeXOS 11.0.

This command was available on the backup MSM in ExtremeXOS 11.0.

The `internal-memory` option was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

mv

```
mv [internal-memory <old-name-internal> internal-memory <new-name-internal>
| internal-memory <old-name-internal> memorycard <new-name-memorycard> |
memorycard <old-name-memorycard> memorycard <new-name-memorycard> |
memorycard <new-name-memorycard> <new-name> | <old-name> memorycard <new-
name-memorycard> | <old-name> <new-name>]
```

Description

Moves or renames an existing configuration, policy, or if configured, core dump file in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
old-name-internal	Specifies the current name of the core dump file located on the internal memory card.
new-name-internal	Specifies the new name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card. NOTE: This parameter is available only on modular switches.
old-name-memorycard	Specifies the current name of the file located on the external compact flash memory card. Depending on your switch configuration, you can have configuration, policy, or cord dump files stored in this card. NOTE: This parameter is available only on modular switches.
new-name-memorycard	Specifies the new name of the file located on the external compact flash memory card. NOTE: This parameter is available only on modular switches.
old-name	Specifies the current name of the configuration or policy file on the system.
new-name	Specifies the new name of the configuration or policy file on the system.

Default

N/A.

Usage Guidelines

If you rename a file with a given extension, remember the following:

- XML-formatted configuration files have the .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have the .xsf file extensions. See [Appendix A, “Software Upgrade and Boot Options,”](#) in the *ExtremeXOS Concepts Guide* for more information.
- Policy files have the .pol file extension.
- Core dump files have the .gz file extension. See [“Internal Memory and Core Dump Files” on page 283](#) for more information.

Make sure the renamed file uses the same file extension as the original file. If you change the file extensions, the file may be unrecognized by the system. For example, if you have an existing configuration file named *test.cfg*, the new filename must include the *.cfg* file extension.

You cannot rename an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to rename the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot rename current selected active configuration file.
```

When you rename a file, the switch displays a message similar to the following:

```
Rename config test.cfg to config megtest.cfg on switch? (y/n)
```

Enter y to rename the file on your system. Enter n to cancel this process and keep the existing filename.

Case-sensitive Filenames. Filenames are case-sensitive. In this example, you have a configuration file named *Test.cfg*. If you attempt to rename the file with the incorrect case, for example *test.cfg*, the switch displays a message similar to the following:

```
Error: mv: unable to rename `/config/test.cfg': No such file or directory
```

Since the switch is unable to locate *test.cfg*, the file is not renamed.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local or remote file, remember the requirements listed above.

Internal Memory and Core Dump Files. Core dump files have a *.gz* file extension. The filename format is: *core.<process-name.pid>.gz* where *process-name* indicates the name of the process that failed and *pid* is the numerical identifier of that process. If you have a modular switch and save core dump files to the external memory card, the filename also includes the affected MSM: MSM-A or MSM-B.

If you configure the switch to send core dump (debug) information to the internal memory card, specify the *internal-memory* option to rename an existing core dump file. If you have a modular switch with an external compact flash memory card installed, you can move and rename the core dump file to that card.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

Modular Switches Only. This command also replicates the action from the primary MSM to the backup MSM. For example, if you rename a file on the primary MSM, the same file on the backup MSM is renamed.

For the `memorycard` option, this command moves files between the external memory card and the switch. If you use the `memorycard` option for both the `old-name` and the `new-name`, this command just renames a file on the external memory card.

For information about core dump files, see the previous section “[Internal Memory and Core Dump Files](#).”

BlackDiamond 10808 Switch Only. In ExtremeXOS 10.1, this command was available only on the primary MSM. For example, if you renamed a file on the primary MSM, the same file on the backup MSM was not renamed.

Example

The following command renames the configuration file named *Testb91.cfg* to *Activeb91.cfg*:

```
mv Testb91.cfg Activeb91.cfg
```

On a modular switch with an external memory card installed, the following command moves the configuration file named *test1.cfg* from the switch to the external memory card:

```
mv test1.cfg memorycard test1.cfg
```

On a modular switch, if you do not change the name of the configuration file, you can also use the following command to move the configuration file *test1.cfg* from the switch to the external memory card:

```
mv test1.cfg memorycard
```

On a modular switch with an external memory card installed, the following command moves the policy file named *bgp.pol* from the memorycard to the switch:

```
mv memorycard bgp.pol bgp.pol
```

History

This command was first available in ExtremeXOS 10.1.

Support for replicating information from the primary MSM to the backup MSM was introduced in ExtremeXOS 11.0.

The `memorycard` option was added in ExtremeXOS 11.1.

The `internal-memory` option was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

restart process

```
restart process [class <cname> | <name> {msm <slot>}]
```

Description

Terminates and restarts the specified process during a software upgrade on the switch.

Syntax Description

cname	Specifies the name of the process to restart. With this parameter, you can terminate and restart all instances of the process associated with a specific routing protocol on all VRs. You can restart the OSPF routing protocol and associated processes.
name	Specifies the name of the process to terminate and restart. You can use this command with the following processes: • bgp • dot1ag (only available on the BlackDiamond 10808 or 12800 series switches) • eaps • exsshd (only available if you have installed the SSH module) • lldp • mpls (only available on the BlackDiamond 10808 or 12800 series switches with MSM1-XLs installed; also requires an MPLS feature pack) • netLogin • ospf • snmpMaster • snmpSubagent • telnetd • tftpd • vrrp
slot	Specifies the MSM where the process should be terminated and restarted. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Use this command to terminate and restart a process during a software upgrade on the switch. You have the following options:

- **cname**—Specifies that the software terminates and restarts all instances of the process associated with a specific routing protocol on all VRs.
- **name**—Specifies the name of the process.

Depending on the software version running on your switch and the type of switch you have, you can terminate and restart different or additional processes. To see which processes you can restart during a software upgrade, enter `restart process` followed by [Tab]. The switch displays a list of available processes.

SummitStack Only. You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: This command can only be executed on Master.
```

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {<name>} {detail} {description} {slot <slotid>}` command. The following is a truncated sample of the `show process` command on the Summit X450 series switch:

Process Name	Version	Restart	State	Start Time
aaa	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
acl	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
bgp	Not Started	0	No license	Not Started
cfgmgr	3.0.0.21	0	Ready	Thu Sep 1 17:00:52 2005
cli	3.0.0.22	0	Ready	Thu Sep 1 17:00:52 2005
devmgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
dirser	3.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
dosprotect	3.0.0.1	0	Ready	Thu Sep 1 17:00:56 2005
eaps	3.0.0.8	0	Ready	Thu Sep 1 17:00:53 2005
...				

You can also use the `restart process` command when upgrading a software modular package. For more information, see the section “[Upgrading a Modular Software Package](#)” in [Appendix A, “Software Upgrade and Boot Options,”](#) of the *ExtremeXOS Concepts Guide*.

Example

The following command stops and restarts the process `tftpd` during a software upgrade:

```
restart process tftpd
```

The following command stops and restarts all instances of the OSPF routing protocol for all VRs during a software upgrade:

```
restart process class ospf
```

History

This command was first available in ExtremeXOS 11.3.

Support for Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) was added in ExtremeXOS 11.4.

Support for MultiProtocol Label Switching (mpls) and Virtual Router Redundancy Protocol (vrrp) was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

rm

```
rm {internal-memory | memorycard} <file-name>
```

Description

Removes/deletes an existing configuration, policy, or if configured, core dump file from the system.

Syntax Description

internal-memory	Specifies the internal memory card.
memorycard	Specifies the removable external compact flash card. NOTE: This parameter is available only on modular switches.
file-name	Specifies the name of the configuration, policy file, or if configured, the core dump file.

Default

N/A.

Usage Guidelines

After you remove a configuration or policy file from the system, that file is unavailable to the system. For information about core dump files, see [“Internal Memory Card and Core Dump Files” on page 289](#).

You cannot remove an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to remove the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot remove current selected active configuration file.
```

When you delete a file from the switch, a message similar to the following appears:

```
Remove testpolicy.pol from switch? (y/n)
```

Enter `y` to remove the file from your system. Enter `n` to cancel the process and keep the file on your system.

Case-sensitive Filenames. Filenames are case-sensitive. In this example, you have a configuration file named `Test.cfg`. If you attempt to remove a file with the incorrect case, for example `test.cfg`, the system is unable to remove the file. The switch does not display an error message; however, the `ls` command continues to display the file `Test.cfg`. To remove the file, make sure you use the appropriate case.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local or remote file, remember the requirements listed above.

Internal Memory Card and Core Dump Files. If you delete a core dump file from the system, that file is unavailable.

If you configure the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to remove/delete the specified core dump file.

For information about configuring and sending core dump information to the internal memory card, see the [configure debug core-dumps](#) and `save debug tracefiles memorycard` commands.

You can use the `*` wildcard to delete core dump files from the internal memory card.

If you configure the switch to write core dump files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if you have a modular switch with an external memory card installed with space available, transfer the files to the external memory card. On the Summit family of switches, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Modular Switches Only. This command also replicates the action from the primary MSM to the backup MSM. For example, if you delete a file on the primary MSM, the same file on the backup MSM is deleted.

For the `memorycard` option, this command removes/deletes an existing file on the card, including core dump files if configured. See the section [“Internal Memory Card and Core Dump Files”](#) for information about core dump files.

You can use the `*` wildcard to delete all of a particular file type from the external memory card; currently running and in use files are not deleted.

BlackDiamond 10808 Switch Only. In ExtremeXOS 10.1, this command was available only on the primary MSM. For example, if you deleted a file on the primary MSM, the same file on the backup MSM was not deleted.

Example

The following command removes the configuration file named *Activeb91.cfg* from the system:

```
rm Activeb91.cfg
```

The following command removes all of the core dump files stored on the internal memory card:

```
rm internal-memory *
```

On a modular switch with an external memory card installed, the following command removes the policy file named `test.pol` from the external memory card:

```
rm memorycard test.pol
```

On a modular switch with an external memory card installed, the following command removes all of the configuration files from the external memory card:

```
rm memorycard *.cfg
```

History

This command was first available in ExtremeXOS 10.1.

Support for replicating information from the primary MSM to the backup MSM was introduced in ExtremeXOS 11.0.

The `memorycard` option was added in ExtremeXOS 11.1.

The `internal-memory` option was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show cpu-monitoring

```
show cpu-monitoring {process <name>} {slot <slotid>}
```

Description

Displays the CPU utilization history of one or more processes.

Command Syntax

name	Specifies the name of the process.
slotid	Specifies the slot number of the MSM module: • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

By default, CPU monitoring is enabled and occurs every 20 seconds. The default CPU threshold value is 60%.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different CPU and process information might be displayed.

SummitStack Only. If you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays CPU utilization history for all the processes running on the master node and the backup node in the Active Topology.
- From the stack manager or a standby node, the stack displays CPU utilization history for all the processes running on the master node and the standby node in the Active Topology.

Modular Switches Only. If you issue the command without any parameters, the switch displays CPU utilization history for all of the processes running on the MSMs installed in your system.

Reading the Output. The `show cpu-monitoring` command is helpful for understanding the behavior of a process over an extended period of time. The following information appears in a tabular format:

- Card—The location (MSM A or MSM B) where the process is running on a modular switch.
- Process—The name of the process.

- Range of time (5 seconds, 10 seconds, and so forth)—The CPU utilization history of the process or the system. The CPU utilization history goes back only 1 hour.
- Total User/System CPU Usage—The amount of time recorded in seconds that the process spends occupying CPU resources. The values are cumulative meaning that the values are displayed as long as the system is running. You can use this information for debugging purposes to see where the process spends the most amount of time: physical memory or virtual memory.

Example

The following command displays CPU utilization on the switch:

```
show cpu-monitoring
```

The following is sample truncated output from a modular switch:

CPU Utilization Statistics - Monitored every 5 seconds

Card	Process	5	10	30	1	5	30	1	Max	Total	
		secs	secs	secs	min	mins	mins	hour		User/System	
		util	util	util	util	util	util	util	util	CPU Usage	
		(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)	(secs)	
MSM-A	System	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.9		
MSM-B	System	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
MSM-A	GNSS_cpuif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_ctrlif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_esmi	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_fabric	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_mac_10g	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pbusmux	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pktengine	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pktif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_switch	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	aaa	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.4	0.82	0.56
MSM-A	acl	0.0	0.0	0.0	0.0	0.0	0.0	0.0	7.5	0.37	0.33
MSM-A	bgp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.2	0.27	0.42
MSM-A	cfgmgr	0.0	0.9	0.3	3.7	1.2	1.2	1.3	27.3	7.70	7.84
MSM-A	cli	0.0	0.0	0.0	48.3	9.6	2.5	2.1	48.3	0.51	0.37
MSM-A	devmgr	0.0	0.0	0.0	0.9	0.3	0.2	0.2	17.1	2.22	2.50
MSM-A	dirser	0.0	0.0	0.0	0.0	0.0	0.0	0.0	9.5	0.0	0.0
MSM-A	dosprotect	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.8	0.20	0.26
MSM-A	eaps	1.9	0.9	0.4	0.0	0.0	0.0	0.0	8.4	2.40	1.40
MSM-A	edp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10.2	0.99	0.47
MSM-A	elrp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.4	0.44	0.28
MSM-A	ems	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.2	1.1	1.16
MSM-A	epm	0.0	0.0	0.0	0.9	0.1	0.2	0.2	4.7	2.6	4.18
MSM-A	esrp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	7.5	0.44	0.36
MSM-A	etmon	0.9	0.4	0.6	1.2	1.1	1.0	1.0	23.3	21.84	7.24
...											

The following is sample output from a Summit X450 series switch:

CPU Utilization Statistics - Monitored every 25 seconds

Process	5	10	30	1	5	30	1	Max	Total	
	secs	secs	secs	min	mins	mins	hour		User/System	
	util	util	util	util	util	util	util	util	CPU Usage	
	(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)	(secs)	
System	n/a	n/a	0.0	0.9	0.1	0.2	0.5	34.6		
aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.72	0.78
acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40	0.24
bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	12.6	11.18	2.21
cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.8	39.8	4743.92	3575.79
cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.59	0.42
devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	19.5	74.44	24.52
dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.8	0.12
eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.1	5.5	36.40	15.41
edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	11.1	10.92	3.97
elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.49	0.44
ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.19	1.29
epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	30.7	48.74	32.93
esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	2.7	0.82	0.45
etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.5	30.5	4865.78	873.87
...										

The following is sample truncated output from a stack:

Slot-1 stackK.3 # sh cpu-monitoring

CPU Utilization Statistics - Monitored every 20 seconds

Card	Process	5	10	30	1	5	30	1	Max	Total	
		secs	secs	secs	min	mins	mins	hour		User/System	
		util	util	util	util	util	util	util	util	CPU Usage	
		(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)	(secs)	
Slot-1	System	n/a	n/a	0.0	1.6	0.8	0.5	0.5	2.5		
Slot-6	System	n/a	n/a	0.3	0.9	0.7	0.4	0.5	4.6		
Slot-1	aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	3.6	1.22	0.75
Slot-1	acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.8	0.52
Slot-1	bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-1	brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.53	0.17
Slot-1	cfgmgr	n/a	n/a	0.1	0.0	0.0	0.0	0.0	0.8	3.18	0.65
Slot-1	cli	n/a	n/a	0.9	0.8	0.1	0.0	0.3	97.2	13.7	2.12
Slot-1	devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.0	1.1	1.24
Slot-1	dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.9	0.0	0.0
Slot-1	dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18	0.12
Slot-1	eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.92	0.45
Slot-1	edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.68	0.20
Slot-1	elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.49	0.21

Slot-1 elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.38	0.34
Slot-1 ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.0	0.41
Slot-1 epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.63	1.28
Slot-1 esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.50	0.21
Slot-1 etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	4.0	0.65
...										
...										
Slot-1 stp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.67	0.27
Slot-1 telnetd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.5	0.23	0.6
Slot-1 tftpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.19	0.10
Slot-1 thttpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.8	0.21	0.13
Slot-1 upm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.43	0.22
Slot-1 vlan	n/a	n/a	0.0	0.0	0.0	0.0	0.1	4.3	4.28	1.56
Slot-1 vrrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.38	0.13
Slot-1 xmld	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.48	0.25
Slot-6 aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.42	0.26
Slot-6 acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40	0.26
Slot-6 bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-6 brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18	0.7
Slot-6 cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.81	0.28
Slot-6 cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	7.17	1.2
Slot-6 devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.35	0.88
Slot-6 dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-6 dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.2
Slot-6 eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.60	0.20
Slot-6 edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.23	0.11
Slot-6 elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.20	0.4
Slot-6 elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.21	0.9
Slot-6 ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.44	0.22
Slot-6 epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.78	1.29
Slot-6 esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.24	0.8
Slot-6 etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.11	0.28
...										

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show heartbeat process

```
show heartbeat process {<name>}
```

Description

Displays the health of the ExtremeXOS processes.

Command Syntax

name	Specifies the name of the process.

Default

N/A.

Usage Guidelines

The software monitors all of the XOS processes running on the switch. This process monitor creates and terminates XOS processes on demand (for example, when you log in or log out of the switch) and restarts processes if an abnormal termination occurs (for example, if your system crashes). The process monitor also ensures that only version-compatible processes and processes with proper licenses are started.

The `show heartbeat process` command is a resource for providing background system health information because you can view the health of ExtremeXOS processes on the switch.

Use this command to monitor the health of the XOS processes. The switch uses two algorithms to collect process health information: polling and reporting. Both polling and reporting measure the heartbeat of the process. Polling occurs when a HELLO message is sent and a HELLO_ACK message is received. The two counts are the same. Reporting occurs when a HELLO_ACK message is sent only. Therefore, no HELLO messages are sent and the HELLO count remains at zero.

The `show heartbeat process` command displays the following information in a tabular format:

- Card—The name of the module where the process is running (modular switches only).
- Process Name—The name of the process.
- Hello—The number of hello messages sent to the process.
- HelloAck—The number of hello acknowledgement messages received by the process manager.
- Last Heartbeat Time—The timestamp of the last health check received by the process manager. (Unknown specifies kernel modules and they do not participate in heartbeat monitoring.)

This status information may be useful for your technical support representative if you have a network problem.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

Example

To display the health of all processes on your system, use the following command:

```
show heartbeat process
```

The following is sample output from a modular switch:

Card	Process Name	Hello	HelloAck	Last Heartbeat Time
MSM-A	aaa	0	180324	Wed Dec 10 15:06:04 2003
MSM-A	acl	36069	36069	Wed Dec 10 15:05:57 2003
MSM-A	bgp	0	180348	Wed Dec 10 15:06:05 2003
MSM-A	cfgmgr	72139	72139	Wed Dec 10 15:06:02 2003
MSM-A	cli	60116	60116	Wed Dec 10 15:06:03 2003
MSM-A	devmgr	0	180339	Wed Dec 10 15:06:03 2003
MSM-A	dirser	0	180324	Wed Dec 10 15:06:03 2003
MSM-A	edp	36069	36069	Wed Dec 10 15:05:57 2003
MSM-A	ems	45087	45087	Wed Dec 10 15:06:03 2003
MSM-A	epm	0	0	Unknown
MSM-A	exacl	0	0	Unknown
....				

The following is sample output from a SummitX450 series switch:

Process Name	Hello	HelloAck	Last Heartbeat Time
aaa	0	254328	Tue Feb 10 05:21:46 2004
acl	50867	50867	Tue Feb 10 05:21:43 2004
bgp	0	0	Wed Feb 4 08:03:18 2004
cfgmgr	25433	25433	Tue Feb 10 05:21:33 2004
cli	84779	84779	Tue Feb 10 05:21:47 2004
cna	20234	20234	Mon Feb 9 00:28:35 2004
devmgr	0	250507	Tue Feb 10 05:21:47 2004
dirser	0	254336	Wed Feb 4 08:03:18 2004
dosprotect	0	254335	Tue Feb 10 05:21:47 2004
eaps	0	254336	Tue Feb 10 05:21:48 2004
edp	50867	50867	Tue Feb 10 05:21:44 2004
elrp	50867	50867	Tue Feb 10 05:21:43 2004
ems	63584	63584	Tue Feb 10 05:21:44 2004
epm	0	0	Wed Feb 4 08:03:18 2004
esrp	50867	50867	Tue Feb 10 05:21:46 2004
...			

To display the health of the STP process on your system, use the following command:

```
show heartbeat process stp
```

The following is sample output from a modular switch:

Card	Process Name	Hello	HelloAck	Last Heartbeat Time
MSM-A	stp	34921	34921	Wed Dec 10 11:54:37 2003

The following is sample output from the Summit X450 series switch:

Process Name	Hello	HelloAck	Last Heartbeat Time
stp	50870	50870	Tue Feb 10 05:22:13 2004

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show memory

```
show memory {slot [slotid | a | b]}
```

Description

Displays the current system memory information.

Syntax Description

slot a	Specifies the MSM module installed in slot A. NOTE: This parameter is available only on modular switches.
slot b	Specifies the MSM module installed in slot B. NOTE: This parameter is available only on modular switches.
slotid	Specifies slot number for the node in a stack. The value can be from 1 to 8.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory process <name> {slot <slotid>}` command to view the system memory and the memory used by the individual processes.

SummitStack Only. If you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays the current system memory information for the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the current system memory information for the master node and the standby node in the Active Topology.

Modular Switches Only. If you issue the command without any parameters, the switch displays information about all of the MSMs installed in your system.

Reading the Output. The `show memory` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM (modular switches only).
- The name of the process.

In general, the free memory count for an MSM or Summit family switch decreases when one or more running processes experiences an increase in memory usage.

If you observe a continuous decrease in the free memory over an extended period of time, and you have not altered your switch configuration, please contact Extreme Networks Technical Support.

Example

The following command displays current system memory information for the MSM installed in slot A of a modular switch:

```
show memory slot a
```

The following is sample output from this command:

System Memory Information

```
MSM-A    Total DRAM (KB): 524288
MSM-A    System      (KB): 45912
MSM-A    User        (KB): 102264
MSM-A    Free        (KB): 376112
```

Memory Utilization Statistics

Card	Slot	Process Name	Memory (KB)
MSM-A	9	aaa	7772
MSM-A	9	acl	6716
MSM-A	9	bgp	16708
MSM-A	9	cfgmgr	3484
MSM-A	9	cli	33964
MSM-A	9	devmgr	3656
MSM-A	9	dirser	3072
MSM-A	9	eaps	9136
MSM-A	9	edp	4644
MSM-A	9	elrp	4608
MSM-A	9	ems	5832
MSM-A	9	epm	8084
MSM-A	9	esrp	11004
MSM-A	9	etmon	11356
MSM-A	9	exacl	13
MSM-A	9	exosmc	22
MSM-A	9	exosq	29
MSM-A	9	exsflow	8
MSM-A	9	exsnoop	15
MSM-A	9	exvlan	252
MSM-A	9	fdb	8760
MSM-A	9	hal	22624
MSM-A	9	mcmgr	13128
MSM-A	9	msgsrv	2972

MSM-A	9	netLogin	4564
MSM-A	9	netTools	4696
MSM-A	9	nettx	56
MSM-A	9	nodemgr	5388
MSM-A	9	ospf	12476
MSM-A	9	pim	10012
MSM-A	9	polMgr	3272
MSM-A	9	rip	10392
MSM-A	9	rtmgr	9748
MSM-A	9	snmpMaster	6400
MSM-A	9	snmpSubagent	8104
MSM-A	9	stp	6896
MSM-A	9	telnetd	3236
MSM-A	9	tftpd	3080
MSM-A	9	vlan	5816
MSM-A	9	vrrp	6584

The following command displays current system memory information for the Summit family of switches:

```
show memory
```

The following is sample output from this command:

System Memory Information

```
-----
Total DRAM (KB): 262144
System      (KB): 25852
User        (KB): 96608
Free        (KB): 139684
```

Memory Utilization Statistics

```
-----
Process Name      Memory (KB)
-----
aaa                13468
acl                11420
bgp                0
cfgmgr            8336
cli               41040
cna                0
devmgr            8452
dirser            7068
dosprotect        8264
eaps              18784
edp               9780
elrp              10040
ems               10672
epm               15520
esrp              16728
etmon             18924
exacl              30
exdos              8
exfib              3
exosmc            29
```


exosnvram	4
exosq	36
exsflow	10
exsnoop	20
exsshd	9272
exvlan	290
fdb	12908
hal	64768
lldp	8816
mcmgr	17836
msgsrv	6960
netLogin	8924
netTools	11524
nettx	70
nodemgr	9636
ospf	18124
ospfv3	0
pim	15996
poe	8936
polMgr	7576
rip	17736
ripng	0
rtmgr	16016
snmpMaster	15416
snmpSubagent	26428
stp	10768
telnetd	8464
tftpd	7584
thttpd	11344
vlan	9660
vrrp	11184
xmld	9148

The following command displays current system memory information for a stack, where slot 1 is the master and slot 6 is the backup:

```
Slot-1 stacK.3 # show memory
```

System Memory Information

```
-----
Slot-1   Total DRAM (KB): 262144
Slot-1   System      (KB): 25476
Slot-1   User        (KB): 132256
Slot-1   Free        (KB): 104412
Slot-6   Total DRAM (KB): 262144
Slot-6   System      (KB): 25476
Slot-6   User        (KB): 122820
Slot-6   Free        (KB): 113848
```

Memory Utilization Statistics

```
-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1   aaa              2548
Slot-1 1   acl              2960
```

Slot-1	1	bgp	0
Slot-1	1	brm	2428
Slot-1	1	cfgmgr	3256
Slot-1	1	cli	16932
Slot-1	1	devmgr	2708
Slot-1	1	dirser	1916
Slot-1	1	dosprotect	1972
Slot-1	1	eaps	6976
Slot-1	1	edp	2656
Slot-1	1	elrp	2640
Slot-1	1	elsm	2592
Slot-1	1	ems	2764
Slot-1	1	epm	3092
Slot-1	1	esrp	2844
Slot-1	1	etmon	16264
...			
Slot-6	6	aaa	2440
Slot-6	6	acl	2872
Slot-6	6	bgp	0
Slot-6	6	brm	2396
Slot-6	6	cfgmgr	2776
Slot-6	6	cli	16292
Slot-6	6	devmgr	2672
Slot-6	6	dirser	1836
Slot-6	6	dosprotect	1944
Slot-6	6	eaps	6924
Slot-6	6	edp	2624
Slot-6	6	elrp	2628
Slot-6	6	elsm	2564
Slot-6	6	ems	2744
Slot-6	6	epm	2976
Slot-6	6	esrp	2792
Slot-6	6	etmon	10068
...			

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show memory process

```
show memory process <name> {slot <slotid>}
```

Description

Displays the current system memory and that of the specified process.

Command Syntax

name	Specifies the name of the process.
slotid	In a modular switch: Specifies the slot number of the MSM module: • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. In a SummitStack, slotid specifies the slot number of the node in the stack topology. The value can be from 1 to 8.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory {slot [slotid | a | b]}` command to view the system memory and the memory used by the individual processes, even for all processes on all MSMs installed in modular switches.

SummitStack Only. If you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays current system memory and that of the specified process running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays current system memory and that of the specified process running on the master node and the standby node in the Active Topology.

Reading the Output. The `show memory process` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM (modular switches only).
- The name of the process.

Example

The following command displays system memory and VRRP memory usage:

```
show memory process vrrp
```

The following is sample output from a modular switch:

```
System Memory Information
-----
MSM-A      Total (KB): 512508 KB
MSM-A      Free  (KB): 395796 KB

Memory Utilization Statistics
-----

Card Slot Process Name      Memory (KB)
-----
MSM-A  9    vrrp              6596
```

The following is sample output from a Summit X450 series switch:

```
System Memory Information
-----
Total DRAM (KB): 262144
System      (KB): 25852
User        (KB): 96608
Free        (KB): 139684

Memory Utilization Statistics
-----

Process Name      Memory (KB)
-----
vrrp              11184
```

The following is sample output from a SummitStack:

```
Slot-1 stack.4 # show memory process "aaa"
```

```
System Memory Information
-----
Slot-1      Total DRAM (KB): 262144
Slot-1      System      (KB): 25476
Slot-1      User        (KB): 132276
Slot-1      Free        (KB): 104392
Slot-6      Total DRAM (KB): 262144
Slot-6      System      (KB): 25476
Slot-6      User        (KB): 122820
Slot-6      Free        (KB): 113848
```

Memory Utilization Statistics

```
-----
```

Card Slot	Process Name	Memory (KB)

Slot-1 1	aaa	2548
Slot-6 6	aaa	2440

History

This command was first available in an ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show process

```
show process {<name>} {detail} {description} {slot <slotid>}
```

Description

Displays the status of the ExtremeXOS processes.

Command Syntax

name	Specifies the name of the process.
detail	Specifies more detailed process information.
description	Describes the name of all of the processes or the specified process running on the switch.
slotid	Specifies the slot number of the MSM module: • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches. In a Summit stack, slotid specifies the slot number of a node in a stack topology.

Default

N/A.

Usage Guidelines

The ExtremeXOS process manager monitors all of the XOS processes. The process manager also ensures that only version-compatible processes are started.

Using this command without the optional keywords displays summary process information. If you specify the `slot` keyword, summary information is displayed for that particular slot only.

The `show process` and `show process slot <slotid>` commands display the following information in a tabular format:

- Card—The name of the module where the process is running (modular switches only).
- Process Name—The name of the process.
- Version—The version number of the process. Options are:
 - Version number—A series of numbers that identify the version number of the process. This is helpful to ensure that you have version-compatible processes and if you experience a problem.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.
- Restart—The number of times the process has been restarted. This number increments by one each time a process stops and restarts.

- **State**—The current state of the process. Options are:
 - **No License**—The process requires a license level that you do not have. For example, you have not upgraded to that license, or the license is not available for your platform.
 - **Ready**—The process is running.
 - **Stopped**—The process has been stopped.
- **Start Time**—The current start time of the process. Options are:
 - **Day/Month/Date/Time/Year**—The date and time the process began. If a process terminates and restarts, the start time is also updated.
 - **Not Started**—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.

If you specify the `detail` keyword, more specific and detailed process information is displayed. The `show process detail` and `show process slot <slotid> detail` commands display the following information in a multi-tabular format:

- Detailed process information
- Memory usage configurations
- Recovery policies
- Process statistics
- Resource usage

This status information may be useful for your technical support representative if you have a network problem.

Depending on the software version running on your switch or your switch model, additional or different process information might be displayed.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

SummitStack Only. If you run the command with out any parameters:

- From the stack manager or backup node, the stack displays the status of the ExtremeXOS processes running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the status of the ExtremeXOS processes running on the standby node and the master node in the Active Topology.

Example

To display the processes on your system, use the following command:

```
show process
```

The following is sample output from a modular switch:

Card	Process Name	Version	Restart	State	Start Time
MSM-A	aaa	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	acl	3.0.0.2	0	Ready	Sat Dec 6 10:54:25 2003
MSM-A	bgp	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	cfgmgr	3.0.0.20	0	Ready	Sat Dec 6 10:54:23 2003

```

MSM-A cli          3.0.0.21      0    Ready    Sat Dec  6 10:54:23 2003
MSM-A devmgr       3.0.0.2       0    Ready    Sat Dec  6 10:54:23 2003
MSM-A dirser       3.0.0.2       0    Ready    Sat Dec  6 10:54:21 2003
MSM-A edp          3.0.0.2       0    Ready    Sat Dec  6 10:54:24 2003
MSM-A ems          3.0.0.2       0    Ready    Sat Dec  6 10:54:23 2003
MSM-A epm          3.0.0.2       0    Ready    Sat Dec  6 10:54:21 2003
MSM-A exacl        3.0.0.2       0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exosmc       3.0.0.2       0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exosq        3.0.0.2       0    Ready    Sat Dec  6 10:54:22 2003
MSM-A exsnoop      3.0.0.2       0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exvlan       3.0.0.2       0    Ready    Sat Dec  6 10:54:22 2003
MSM-A fdb          3.0.0.2       0    Ready    Sat Dec  6 10:54:24 2003
....

```

The following is sample output from a Summit X450 series switch:

Process Name	Version	Restart	State	Start Time
aaa	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
acl	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
bgp	Not Started	0	No license	Not Started
cfgmgr	3.0.0.21	0	Ready	Thu Sep 1 17:00:52 2005
cli	3.0.0.22	0	Ready	Thu Sep 1 17:00:52 2005
devmgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
dirser	3.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
dosprotect	3.0.0.1	0	Ready	Thu Sep 1 17:00:56 2005
eaps	3.0.0.8	0	Ready	Thu Sep 1 17:00:53 2005
edp	3.0.0.2	0	Ready	Thu Sep 1 17:00:53 2005
elrp	3.0.0.1	0	Ready	Thu Sep 1 17:00:53 2005
ems	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
epm	3.0.0.4	0	Ready	Thu Sep 1 17:00:49 2005
esrp	3.0.0.4	0	Ready	Thu Sep 1 17:00:53 2005
etmon	1.0.0.1	0	Ready	Thu Sep 1 17:00:55 2005
exacl	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exdos	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exfib	1.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
exosmc	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exosnvram	3.0.0.3	0	Ready	Thu Sep 1 17:00:50 2005
exosq	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exsflow	1.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
exsnoop	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exvlan	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
fdb	4.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
hal	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
lacp	3.0.0.1	0	Ready	Thu Sep 1 17:00:53 2005
lldp	1.2.0.0	0	Ready	Thu Sep 1 17:00:53 2005
mcmgr	4.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
msgsrv	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
netLogin	1.0.0.0	0	Ready	Thu Sep 1 17:00:55 2005
netTools	3.0.0.2	0	Ready	Thu Sep 1 17:00:55 2005
nettx	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
nodemgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
ospf	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
ospfv3	Not Started	0	No license	Not Started
pim	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
poe	3.0.0.2	0	Ready	Thu Sep 1 17:00:56 2005


```

polMgr      3.0.0.2    0    Ready    Thu Sep 1 17:00:54 2005
rip         3.0.0.2    0    Ready    Thu Sep 1 17:00:54 2005
ripng       3.0.0.1    0    Ready    Thu Sep 1 17:00:54 2005
rtmgr       4.0.0.2    0    Ready    Thu Sep 1 17:00:53 2005
snmpMaster  3.0.0.2    0    Ready    Thu Sep 1 17:00:52 2005
snmpSubagent 3.0.0.2    0    Ready    Thu Sep 1 17:00:52 2005
stp         3.0.3.1    0    Ready    Thu Sep 1 17:00:53 2005
telnetd     3.0.0.2    0    Ready    Thu Sep 1 17:00:55 2005
tftpd       3.0.0.2    0    Ready    Thu Sep 1 17:00:55 2005
thttpd      3.0.0.1    0    Ready    Thu Sep 1 17:00:55 2005
vlan        3.1.0.2    0    Ready    Thu Sep 1 17:00:52 2005
vrrp        3.0.0.5    0    Ready    Thu Sep 1 17:00:55 2005
xmld        1.0.0.0    0    Ready    Thu Sep 1 17:00:56 2005

```

The following example specifies the process aaa along with the detail keyword:

```
show process aaa detail
```

The following is sample output from this command:

```

Name          PID      Path    Type Link Date                               Build By      Peer
-----
aaa           284     ./aaa   App  Thu Dec 4 13:23:07 PST 2003  release-manager 2
3
Virtual Router(s):
-----
Configuration:
Start Priority  SchedPolicy  Stack  TTY  CoreSize  Heartbeat  StartSeq
-----
1           0           0       0     0       0           1         1
Memory Usage Configuration:
Memory(KB) Zones: Green Yellow Orange Red
-----
0           0           0       0     0
Recovery policies
-----
failover-reboot
-----
Statistics:
ConnetionLost  Timeout  Start  Restart  Kill  Register  Signal  Hello  Hello Ack
-----
0             0         0       0         0     1         0       0      173199
Memory Zone  Green  Yellow  Orange  Red
-----
Green       0       0       0       0
-----
Commands:
Start      Stop      Resume      Shutdown      Kill
-----
0          0         0           0             0
-----
Resource Usage:
UserTime SysTime  PageReclaim PageFault Up Since          Up Date  Up Time
-----

```

```
2.160000 0.560000      546      966   Sat Dec  6 10:54:24 2003 00/00/04 00:14:02
```

Thread Name	Pid	Tid	Delay	Timeout	Count
tacThread	0	2051	10	0	
radiusThread	0	1026	10	1	
main	0	1024	2	1	

The following example describes the name of all of the processes running on the switch:

```
show process description
```

The following is sample output from this command:

Process Name	Description
aaa	Authentication, Authorization, and Accounting Server
acl	Access Control List Manager
bgp	Border Gateway Protocol
cfgmgr	Configuration Manager
cli	Cli Manager
devmgr	Device Manager
dirser	Directory Services
eaps	Ethernet Automatic Protection Switching
edp	Extreme Discovery Protocol
elrp	Extreme Loop Recognition Protocol
ems	Event Management System Server
epm	Extreme Process Manager
esrp	Extreme Standby Routing Protocol
etmon	Traffic monitoring and sampling utility
exacl	Access Control List Module
exfib	Routing interface to manage missing routes in ASIC
exosmc	Multicast Forwarding Module
exosq	EXOS Queue Module
exsflow	Sflow interface to gather sflow samples
exsnoop	IGMP/MLD Snooping Module
exvlan	Layer 2 configuration module
fdb	Forwarding Data Base Manager
hal	Hardware Abstraction Layer
lldp	802.1AB; Station and Media Access Control Connectivity Discover
mcmgr	Multicast Cache Manager
msgsrv	Message Server
netLogin	Network Login includes MAC, Web-Based and 802.1X authentication
netTools	Network Toolset includes ping/tracert/bootprelay/dhcp/dns/sntp
nettx	Layer 2 forwarding engine module
nodemgr	Fault Tolerance Manager
ospf	Open Shortest Path First Routing Protocol
ospfv3	Open Shortest Path First Routing Protocol for IPv6
pim	Protocol Independent Multicast
polMgr	Policy Manager
rip	Routing Information Protocol
ripng	Routing Information Protocol for IPv6
rtmgr	Route Table Manager
snmpMaster	Simple Network Management Protocol - Master agent
snmpSubagent	Simple Network Management Protocol - Subagent

```

stp          Spanning Tree Protocol
telnetd      Telnet server
tftpd        Tftp server
thttpd       Web Server
vlan         VLAN Manager - L2 Switching application
vrrp         Virtual Router Redundancy Protocol (RFC 3768)
xmld         XML server

```

The following example shows the truncated output for the command on a stack:

Slot-1 stack.7 # show process

Card	Process Name	Version	Restart	State	Start Time
Slot-1	aaa	3.0.0.3	0	Ready	Thu Mar 1 11:29:34 2007
Slot-1	acl	3.0.0.2	0	Ready	Thu Mar 1 11:29:39 2007
Slot-1	bgp	Not Started	0	No license	Not Started
Slot-1	brm	1.0.0.0	0	Ready	Thu Mar 1 11:29:42 2007
Slot-1	cfgmgr	3.0.0.21	0	Ready	Thu Mar 1 11:29:34 2007
Slot-1	cli	3.0.0.22	0	Ready	Thu Mar 1 11:29:34 2007
Slot-1	devmgr	3.0.0.2	0	Ready	Thu Mar 1 11:29:34 2007
Slot-1	dirser	3.0.0.2	0	Ready	Thu Mar 1 11:29:33 2007
Slot-1	dosprotect	3.0.0.1	0	Ready	Thu Mar 1 11:29:41 2007
Slot-1	eaps	3.0.0.8	0	Ready	Thu Mar 1 11:29:36 2007
Slot-1	edp	3.0.0.2	0	Ready	Thu Mar 1 11:29:36 2007
Slot-1	elrp	3.0.0.1	0	Ready	Thu Mar 1 11:29:35 2007
Slot-1	elsm	3.0.0.1	0	Ready	Thu Mar 1 11:29:35 2007
Slot-1	ems	3.0.0.2	0	Ready	Thu Mar 1 11:29:34 2007
Slot-1	epm	3.0.0.4	0	Ready	Thu Mar 1 11:29:30 2007
Slot-1	esrp	3.0.0.4	0	Ready	Thu Mar 1 11:29:37 2007
Slot-1	etmon	1.0.0.1	0	Ready	Thu Mar 1 11:29:40 2007
...					
...					
Slot-6	aaa	3.0.0.3	0	Ready	Thu Mar 1 11:29:22 2007
Slot-6	acl	3.0.0.2	0	Ready	Thu Mar 1 11:29:24 2007
Slot-6	bgp	Not Started	0	No license	Not Started
Slot-6	brm	1.0.0.0	0	Ready	Thu Mar 1 11:29:26 2007
Slot-6	cfgmgr	3.0.0.21	0	Ready	Thu Mar 1 11:29:21 2007
Slot-6	cli	3.0.0.22	0	Ready	Thu Mar 1 11:29:21 2007
Slot-6	devmgr	3.0.0.2	0	Ready	Thu Mar 1 11:29:21 2007
Slot-6	dirser	3.0.0.2	0	Ready	Thu Mar 1 11:29:20 2007
Slot-6	dosprotect	3.0.0.1	0	Ready	Thu Mar 1 11:29:26 2007
Slot-6	eaps	3.0.0.8	0	Ready	Thu Mar 1 11:29:23 2007
Slot-6	edp	3.0.0.2	0	Ready	Thu Mar 1 11:29:23 2007
Slot-6	elrp	3.0.0.1	0	Ready	Thu Mar 1 11:29:22 2007
Slot-6	elsm	3.0.0.1	0	Ready	Thu Mar 1 11:29:22 2007
Slot-6	ems	3.0.0.2	0	Ready	Thu Mar 1 11:29:21 2007
Slot-6	epm	3.0.0.4	0	Ready	Thu Mar 1 11:29:19 2007
Slot-6	esrp	3.0.0.4	0	Ready	Thu Mar 1 11:29:23 2007
Slot-6	etmon	1.0.0.1	0	Ready	Thu Mar 1 11:29:25 2007
...					

History

This command was first available in an ExtremeXOS 10.1.

The `description` keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all available platforms.

start process

```
start process <name> {msm <slot>}
```

Description

Restarts the specified process on the switch.

Syntax Description

name	Specifies the name of the process to start. You can start the following processes: • bgp • eaps • exsshd (only available if you have installed the SSH module) • lldp • netLogin • ospf • snmpMaster • snmpSubagent • telnetd • tftpd • vrrp
slot	Specifies the MSM where the process should be started. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Use this command after you have stopped a process and you want to restart it. To stop a process, use the [terminate process](#) command.

You are unable to start a process that is already running. If you try to start a currently running process, an error message similar to the following appears:

```
Error: Process telnetd already exists!
```

Depending on the software version running on your switch and the type of switch you have, you can restart different or additional processes. To see which processes you can restart, enter `start process` followed by [Tab]. The switch displays a list of available processes.

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {<name>} {detail} {description} {slot <slotid>}` command.

You can also use the `start process` command when upgrading a software modular package. For more information, see the section “[Upgrading a Modular Software Package](#)” in [Appendix A, “Software Upgrade and Boot Options,”](#) of the *ExtremeXOS Concepts Guide*.

SummitStack Only. You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: This command can only be executed on Master.
```

Example

The following restarts the process *tftpd*:

```
start process tftpd
```

History

This command was first available in ExtremeXOS 11.0.

Support for restarting the Link Layer Discovery Protocol (lldp), Open Shortest Path First (ospf), and network login (netLogin) processes was added in ExtremeXOS 11.3.

Support for restarting the Border Gateway Protocol (bgp) was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

terminate process

```
terminate process <name> [forceful | graceful] {msm <slot>}
```

Description

Terminates the specified process on the switch.

Syntax Description

name	Specifies the name of the process to terminate. You can terminate the following processes: • bgp • dot1ag (only available on the BlackDiamond 10808 or 12800 series switches) • eaps • exsshd (only available if you have installed the SSH module) • lldp • mpls (only available on the BlackDiamond 10808 or 12800 series switches with MSM1-XLs installed; also requires an MPLS feature pack) • netLogin • ospf • snmpMaster • snmpSubagent • telnetd • thttpd • tftpd • vrrp
forceful	Specifies a forceful termination.
graceful	Specifies a graceful termination.
slot	For a modular chassis, specifies the MSM where the process should be terminated. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B. On a SummitStack, specifies the target node's slot number. The number is a value from 1 to 8. NOTE: This parameter is available only on modular switches and SummitStack.

Default

N/A.

Usage Guidelines

If recommended by Extreme Networks Technical Support personnel, you can stop a running process.

The `forceful` option quickly terminates a process on demand. Unlike the `graceful` option, the process is immediately shutdown without any of the normal process cleanup. The status of the operation is displayed on the console. After a successful forceful termination of a process, a message similar to the following appears:

```
Forceful termination success for snmpMaster
```

The `graceful` option terminates the process by allowing it to close all opened connections, notify peers on the network, and other types of process cleanup. After this phase, the process is finally terminated. After a successful graceful termination of a process, a message similar to the following appears:

```
Successful graceful termination for snmpSubagent
```

SummitStack Only. You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: This command can only be executed on Master.
```

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {<name>} {detail} {description} {slot <slotid>}` command.

Depending on the software version running on your switch and the type of switch you have, you can terminate different or additional processes. To see which processes you can terminate, enter `terminate process` followed by [Tab]. The switch displays a list of available processes.

To restart a process that has been terminated, use the `start process` command.



NOTE

Do not terminate a process that was installed since the last reboot unless you have saved your configuration. If you have installed a software module and you terminate the newly installed process without saving your configuration, your module may not be loaded when you attempt to restart the process with the `start process` command.

To preserve a process's configuration during a terminate and (re)start cycle, save your switch configuration before terminating the process. Do not save the configuration during the process terminate and re(start) cycle. If you save the configuration after terminating a process, and before the process (re)starts, the configuration for that process is lost.

You can also use the `terminate process` command when upgrading a software modular package. For more information, see the section “[Upgrading a Modular Software Package](#)” in [Appendix A, “Software Upgrade and Boot Options,”](#) of the *ExtremeXOS Concepts Guide*.

Example

The following initiates a graceful termination of the process `tftpd`:

```
terminate process tftpd graceful
```

History

This command was first available in ExtremeXOS 11.0.

Support for terminating the Link Layer Discovery Protocol (lldp), network login (netLogin), and Open Shortest Path First (ospf) processes was added in ExtremeXOS 11.3.

Support for terminating the Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) processes was added in ExtremeXOS 11.4.

Support for terminating the MultiProtocol Label Switch (mpls) and Virtual Router Redundancy Protocol (vrrp) processes was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

5

Commands for Configuring Slots and Ports on a Switch

This chapter describes commands related to:

- Enabling, disabling, and configuring individual ports
- Configuring port speed (Fast Ethernet ports only) and half- or full-duplex mode
- Creating link aggregation groups on multiple ports
- Displaying port statistics
- Configuring mirroring
- Configuring software-controlled redundant ports and Smart Redundancy
- Configuring Extreme Discovery Protocol

By default, all ports on the switch are enabled. After you configure the ports to your specific needs, you can select which ports are enabled or disabled.

Fast Ethernet ports can connect to either 10BASE-T or 100BASE-T networks. By default, the ports autonegotiate (automatically determine) the port speed. You can also configure each port for a particular speed (either 10 Mbps or 100 Mbps).

In general Gigabit Ethernet ports with fiber interfaces are statically set, and their speed cannot be modified. However, there are two GBICs supported by Extreme that can have a configured speed:

- 100 FX GBICs, which must have their speed configured to 100 Mbps
- 100FX/1000LX GBICs, which can be configured at either speed

The switch comes configured to use autonegotiation to determine the port speed and duplex setting for each port. You can select to manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on gigabit Ethernet ports.

All ports on the switch (except gigabit Ethernet ports) can be configured for half-duplex or full-duplex operation. The ports are configured to autonegotiate the duplex setting, but you can manually configure the duplex setting for your specific needs.

Flow control is supported only on gigabit Ethernet ports. It is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. (See the *ExtremeXOS Concepts Guide* for more detailed information on flow control on Extreme Networks devices.)

Link aggregation, or load sharing, with Extreme Network switches allows you to increase bandwidth and resilience between switches by using a group of ports to carry traffic in parallel between switches. The sharing algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. The algorithm also guarantees packet sequencing between clients.

Beginning with ExtremeXOS version 11.3, you can use the Link Aggregation Control Protocol (LACP) to configure and enable dynamic link aggregation. Static link aggregation is still available as well.

If a port in a link aggregation group fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

You can view port status on the switch using the `show ports` commands. These commands, when used with specific keywords and parameters, allow you to view various issues such as collision statistics, link speed, flow control, and packet size. Beginning with ExtremeXOS software version 11.3, these port information displays show real-time statistics, or you can configure the display to show a snapshot of real-time statistics (as in earlier versions of the software).

Beginning with ExtremeXOS version 11.6 software, you can configure WAN PHY OAM on those interfaces that connect 10G Ethernet ports to the SONET/SDH network.

Commands that require you to enter one or more port numbers use the parameter `<port_list>` in the syntax. On a modular switch or SummitStack, a `<port_list>` can be a list of slots and ports. On a stand-alone switch, a `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see [“Port Numbering”](#) in [Chapter 1, “Command Reference Overview.”](#)

clear counters ports

```
clear counters ports
```

Description

Clears the counters associated with the ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines



NOTE

If you use the `clear counters` command with no keyword, the system clears the counters for all applications.

This command clears the counters for the ports, including the following:

- Statistics
- Transmit errors
- Receive errors
- Collisions
- Packets

Example

The following command clears the counters on all ports:

```
clear counters ports
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

clear lacp counters

```
clear lacp counters
```

Description

Clears the counters associated with Link Aggregations Control Protocol (LACP).

Syntax Description

This command has no parameters or variables.

Default

N/A

Usage Guidelines

This command clears the following counters for LACP; it sets these counters back to 0 for every LACP port on the device:

- LACP PDUs dropped on non_LACP ports
- Stats
 - Rx - Accepted
 - Rx - Dropped due to error in verifying PDU
 - Rx - Dropped due to LACP not being up on this port
 - Rx - Dropped due to matching own MAC
 - Tx - Sent Successfully
 - Tx - Transmit error

Example

The following command clears the LACP counters on all ports:

```
clear lacp counters
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

clear counters edp

```
clear counters edp {ports <ports>}
```

Description

Clears the counters associated with Extreme Discovery Protocol (EDP).

Syntax Description

ports	Specifies one or more ports or slots and ports.

Default

If you do not specify a port, the EDP counters will be cleared for all ports.

Usage Guidelines

This command clears the following counters for EDP protocol data units (PDUs) sent and received per EDP port:

- Switch PDUs transmitted
- VLAN PDUs transmitted
- Transmit PDUs with errors
- Switch PDUs received
- VLAN PDUs received
- Received PDUs with errors

Example

The following command clears the EDP counters on all ports:

```
clear counters edp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear slot

```
clear slot <slot>
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

slot	Specifies the slot number.

Default

N/A.

Usage Guidelines

All configuration information related to the slot and the ports on the module is erased. If a module is present when you issue this command, the module is reset to default settings.

If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state (where the inserted module does not match the configured slot), and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type. Use the `configure slot module` command to configure the slot.

Example

The following command clears slot 2 of a previously assigned module type:

```
clear slot 2
```

The following command clears slot 4 of a previously assigned module type in a stack:

```
clear slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on modular switches and SummitStack.

configure edp advertisement-interval

```
configure edp advertisement-interval <timer> holddown-interval <timeout>
```

Description

Sets the advertisement interval and hold down interval for EDP.

Syntax Description

timer	Specifies the advertisement interval in seconds.
timeout	Specifies the hold down interval in seconds.

Default

The default setting for timer is 60 seconds, and for timeout is 180 seconds.

Usage Guidelines

Extreme Discover Protocol (EDP) is used to gather information about neighbor Extreme Networks switches. EDP-enabled ports advertise information about the Extreme switch to other switches on the interface and receive advertisements from other Extreme switches. Information about other Extreme switches is discarded after the hold down interval timeout value is reached without receiving another advertisement.

Example

The following command configures the EDP advertisement-interval to 2 minutes and the hold down interval to 6 minutes:

```
configure edp advertisement-interval 120 holddown-interval 360
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ip-mtu vlan

```
configure ip-mtu <mtu> vlan <vlan_name>
```

Description

Sets the maximum transmission unit (MTU) for the VLAN.

Syntax Description

mtu	Specifies the IP maximum transmission unit (MTU) value. Range is from 1500 to 9216.
vlan_name	Specifies a VLAN name.

Default

The default IP MTU size is 1500.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit series switches do not support fragmentation of any IP packets they forward. Also, the BlackDiamond 8800 series switch, SummitStack, and the Summit series switches do not support the router specification for path MTU discovery.

Beginning with ExtremeXOS 11.6, the BlackDiamond 8800 a-series and e-series switches and the Summit X450a and X450e series switches support IP fragmentation and path MTU discovery; the Summit X250e series switches are included with ExtremeXOS 12.0.

Use this command to enable jumbo frame support or for IP fragmentation with jumbo frames. Jumbo frames are Ethernet frames that are larger than 1522 bytes, including 4 bytes used for CRC. Both endstations involved in the transfer must be capable of supporting jumbo frames. The switch does not perform IP fragmentation or participate in MTU negotiation on behalf of devices that do not support jumbo frames.

When enabling jumbo frames and setting the MTU size for the VLAN, keep in mind that some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC included in a jumbo frame configuration. Ensure that the NIC maximum MTU is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

If you use IP fragmentation with jumbo frames and you want to set the MTU size greater than 1500, all ports in the VLAN must have jumbo frames enabled.

Example

The following command sets the MTU size to 2000 for VLAN *sales*:

```
configure ip-mtu 2000 vlan sales
```

History

This command was available in ExtremeXOS 11.0.

Platform Availability

This command is available on the BlackDiamond 10808, the BlackDiamond 12800 series switches, the BlackDiamond 8800 a-series and e-series modules, and the Summit X450a, X450e and X250e series switches (whether or not included in a SummitStack).

configure jumbo-frame-size

```
configure jumbo-frame-size <framesize>
```

Description

Sets the maximum jumbo frame size for the switch.

Syntax Description

framesize	Specifies a maximum transmission unit (MTU) size for a jumbo frame. The range is 1523 to 9216; the default is 9216.

Default

Jumbo frames are disabled by default. The default size setting is 9216.

Usage Guidelines

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations involved in the transfer must be capable of supporting jumbo frames.

The `framesize` keyword describes the maximum jumbo frame size “on the wire,” and includes 4 bytes of cyclic redundancy check (CRC) plus another 4 bytes if 802.1Q tagging is being used.

To enable jumbo frame support, you must configure the maximum transmission unit (MTU) size of a jumbo frame that will be allowed by the switch.



NOTE

Extreme Networks recommends that you set the MTU size so that fragmentation does not occur.

Some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU size is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

Example

The following command configures the jumbo frame size to 5500:

```
configure jumbo-frame-size 5500
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure lacp member-port priority

```
configure lacp member-port <port> priority <port_priority>
```

Description

Configures the member port of an LACP to ensure the order that ports are added to the aggregator. The lower value you configure for the port's priority, the higher priority that port has to be added to the aggregator.

Syntax Description

port	Specifies the LACP member port that you are specifying the priority for.
port_priority	Specifies the priority you are applying to this member port to be assigned to the LACP aggregator. The range is from 1 to 65535; the default is 0. The lower configured value has higher priority to be added to the aggregator.

Default

The default priority is 0.

Usage Guidelines

The port must be added to the LAG prior to configuring it for LACP. The default value is 0, or highest priority.

You can configure the port priority to ensure the order in which LAG ports join the aggregator. If you do not configure this parameter, the lowest numbered ports in the LAG are the first to be added to the aggregator; if there are additional ports configured for that LAG, they are put in standby mode.

Use this command to override the default behavior and ensure the order in which LAG ports are selected. Also, if more than one port is *configured* with the same priority, the lowest numbered port joins the aggregator.

Example

The following command sets the port priority for the LAG port 5:1 to be 55 (which will probably put that port in standby initially):

```
configure lacp member-port 5:1 priority 55
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mirror add anomaly

Configure mirror add ports <port list> anomaly

Description

Mirrors detected anomaly traffic to the mirror port.

Syntax Description

port list	Specifies the list of ports.

Default

N/A.

Usage Guidelines

The command mirrors detected anomaly traffic to the mirror port. You must enable a mirror port and enable protocol anomaly protection on the slot that has the port to be monitored before using this command. After configuration, only detected anomaly traffic from these ports are dropped or mirrored to the mirror port, and legitimate traffic is not affected.

This command takes effect after enabling anomaly-protection.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure mirroring add

```
configure mirroring add [vlan <name> {port <port>} | port <port> {vlan
<name>}] {ingress | egress | ingress-and-egress}
```

Description

Adds a particular mirroring filter definition on the switch.

Syntax Description

vlan	Specifies a VLAN.
name	Specifies a VLAN name.
port	Specifies a port or slot and port.
port	Specifies particular ports or slots and ports.
ingress	Specifies packets be mirrored as they are received on a port. NOTE: This parameter is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches and only with port-based mirroring.
egress	Specifies packets be mirrored as they are sent from a port. NOTE: This parameter is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches and only with port-based mirroring.
ingress-and-egress	Specifies all forwarded packets be mirrored. This is the default setting on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches for port-based mirroring. NOTE: This parameter is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches and only with port-based mirroring.

Default

N/A.

Usage Guidelines

You must enable port-mirroring using the `enable mirroring to port` command before you can configure the mirroring filter definitions.

Port mirroring configures the switch to copy all traffic associated with one or more ports to a monitor port on the switch. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 16 mirroring filters and one monitor port can be configured on the switch. Frames that contain errors are not mirrored.

Guidelines for configuring mirroring on all switches. This section summarizes the guidelines for configuring mirroring:

- When you disable mirroring, all the filters are unconfigured.
- You cannot mirror the monitor port.
- The mirroring configuration is removed when you:
 - Delete a VLAN (for all VLAN-based filters).
 - Delete a port from a VLAN (for all VLAN-, port-based filters).
 - Unconfigure a slot (for all port-based filters on that slot).
- Any mirrored port can also be enabled for load sharing (or link aggregation); however, each individual port of the load-sharing group must be explicitly configured for mirroring.
- The mirroring filters are not confined to a single module; they can have ports that span multiple modules.
- You cannot use the management port at all in mirroring configurations.
- You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

```
Error: Port mirroring cannot enabled on an ELSM enabled port 1.
```

Summit family of switches only. The traffic filter can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.
 (If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).
- VLAN—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.
- Only 8 VLANs can be mirrored on a given physical port.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- In standard mode (see `configure mirroring mode` command), even if you select ingress and egress traffic, the packet is mirrored only the *first* time it matches a mirror filter and is not mirrored on subsequent configured filters. In enhanced mode, packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- You cannot include the monitor port for the Summit family of switches in a load-sharing group.
- You cannot run sFlow and mirroring on the same switch on Summit X450 series switches. If you attempt to enable mirroring on these devices on a port that is already enabled for sFlow, the switch returns the following message:

```
Mirroring is not compatible with SFlow. Mirroring is not enabled!
```

You can run mirroring and sFlow on the same device when you are running one of the following:

- Summit X450a or X450e or X250e series switches
- Tagged and untagged traffic is mirrored slightly differently depending on the switch that the mirrored port and the monitor port are on:
 - With a monitor port on a Summit X450 series switch (original Summit X450 switch), *all* traffic egressing the monitor port is tagged. Even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged with the internal VLAN ID.
 - With a monitor port on a Summit X450a, X450e or X250e series switch, the mirrored packet is tagged *only* if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port as untagged.
 - Egress mirrored packets are always transmitted as tagged on the monitor port(s).

BlackDiamond 8800 series switch and SummitStack only. The traffic filter can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).
- VLAN—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.
- Only 8 VLANs can be mirrored on a given physical port.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- In standard mode (see `configure mirroring mode` command), even if you select ingress and egress traffic, the packet is mirrored only the *first* time it matches a mirror filter and is not mirrored on subsequent configured filters. In enhanced mode, packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- You cannot include the monitor port for a BlackDiamond 8800 series switch or SummitStack in a load-sharing group.
- You cannot run sFlow and mirroring when using the BlackDiamond 8800 original modules or when using Summit X450 switches in a SummitStack. If you attempt to enable mirroring on these devices on a port that is already enabled for sFlow, the switch returns the following message:

```
Mirroring is not compatible with SFlow. Mirroring is not enabled!
```

You can run mirroring and sFlow on the same device when you are running one of the following:

- BlackDiamond a-series and e-series modules
- Summit X450a, X450e, or X250e switches in a SummitStack

- Tagged and untagged traffic is mirrored slightly differently depending on the module that the mirrored port and the monitor port are on:
 - With a monitor port on a BlackDiamond 8800 original module, *all* traffic egressing the monitor port is tagged (regardless of what module the ingress port is on). Even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged with the internal VLAN ID.
 - With a monitor port on a BlackDiamond 8800 a-series or e-series module, the mirrored packet is tagged *only* if the ingress packet is tagged (regardless of what module the ingress port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port as untagged.
 - Egress mirrored packets are always transmitted as tagged on the monitor port(s).
- With the BlackDiamond 8800 a-series and e-series modules, you may see a packet mirrored twice. This occurs only if both the ingress mirrored port and the monitor port are on the same one-half of the module *and* the egress mirrored port is either on the other one-half of that module or on another module.

BlackDiamond 10808 and BlackDiamond 12800 series switches only. The traffic filter can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.
- VLAN—All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.
- The master port of the load-sharing group can be the monitor port for port-mirroring.

The monitor port transmits tagged or untagged frames, according to the way you configured the monitor port. This feature allows you to mirror multiple ports or VLANs to a monitor port, while preserving the ability of a single protocol analyzer to track and differentiate traffic within a broadcast domain (VLAN) and across broadcast domains (for example, across VLANs when routing).



NOTE

The monitor port on the BlackDiamond 10808 switch must be explicitly configured for tagged or untagged frames beginning with ExtremeXOS version 11.0.

The traffic egressing the monitor port can be either tagged or untagged. If the mirroring is enabled as tagged on the monitor port, all traffic egressing the monitor port is tagged. In this case, even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port as tagged. And, if mirroring is enabled as untagged on the monitor port, all traffic egressing the monitor port is untagged, including mirrored tagged packets.

When you upgrade to ExtremeXOS 11.0 on the BlackDiamond 10808 switches, all restored mirroring configurations are tagged on the monitor ports.

Example

The following example sends all traffic coming into an BlackDiamond 8800 series switch on slot 3, port 2 to the mirror port:

```
configure mirroring add port 3:2 ingress
```

The following example sends all traffic coming into a Summit X450 series switch on port 11 and the VLAN *default* to the mirror port:

```
configure mirroring add port 11 vlan default
```

The following example sends all traffic coming into or out of a BlackDiamond 10808 switch on slot 3, port 2 and the VLAN *default* to the mirror port:

```
configure mirroring add port 3:2 vlan default
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was added in ExtremeXOS 11.0.

The ingress/egress option was added in ExtremeXOS 11.1.

The VLAN mirroring capability was added to the BlackDiamond 8800 series switch in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mirroring delete

```
configure mirroring delete [all | port <port> {vlan <name>} |vlan <name>
{port <port>}]
```

Description

Deletes a particular mirroring filter definition on the switch.

Syntax Description

all	Specifies all mirroring filter definitions.
port	Specifies a port or a slot and port.
port	Specifies particular ports or slots and ports.
vlan	Specifies a VLAN.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

On a modular switch, <port_list> must be a slot and port in the form <slot>:<port>. For a detailed explanation of port specification, see [“Port Numbering”](#) in [Chapter 1, “Command Reference Overview.”](#)

Example

The following example deletes the mirroring filter on an BlackDiamond 8800 series switch defined for slot 7, port 1:

```
configure mirroring delete ports 7:1
```

The following example deletes the mirroring filter on a BlackDiamond 10808 switch defined for slot 3, port 2 on VLAN default:

```
configure mirroring delete ports 3:2 vlan default
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was added in ExtremeXOS 11.0.

The VLAN mirroring capability was added to the BlackDiamond 8800 series switch in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mirroring mode

```
configure mirroring mode [enhanced | standard]
```

Description

Configures the mirroring mode which affects mirroring behavior globally in the system.

Syntax Description

enhanced	Specifies the mirroring mode that provides enhanced mirroring operation.
standard	Specifies the standard mirroring mode that is required when the mirroring configuration involves ports or VLANs on non “a” or “e” series hardware on Summit X450 series switches and BlackDiamond 8800 modules.

Default

Standard mode is the default.

Usage Guidelines

When the mirroring configuration involves only ports of VLANs on “a” or “e” series hardware, enhanced mode is recommended since it provides enhanced behavior. (Refer to “Guidelines for Mirroring on the Summit Family of Switches Only” and “Guidelines for Mirroring on the BlackDiamond 8800 Series Switch Only” in Chapter 5 of the XOS Concepts Guide.)

Example

The following example configures a system to use enhanced mirroring mode:

```
configure mirroring mode enhanced
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on Summit family, SummitStack, and BlackDiamond 8800 series switches.

configure ports auto off

```
configure ports <port_list> auto off speed [10 | 100 | 1000 | 10000] duplex
[half | full]
```

Description

Manually configures port speed and duplex setting configuration on one or more ports on a switch.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
speed [10]	Specifies 10 Mbps ports.
speed [100]	Specifies 100 Mbps ports.
speed [1000]	Specifies 1000 Mbps (1 Gigabit) ports.
speed [10000]	Specifies 10000 Mbps (10 Gigabit) ports.
duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on for 1 G ports.

Auto off for 10 G ports.

Usage Guidelines

You can manually configure the duplex setting and the speed on 10/100 and 10/100/1000 Mbps and fiber SFP gigabit Ethernet ports.

In general, SFP gigabit Ethernet ports are statically set to 1 Gbps, and their speed cannot be modified. However, there are two GBICs supported by Extreme that can have a configured speed:

- 100 FX GBICs, which must have their speed configured to 100 Mbps
- 100FX/1000LX GBICs, which can be configured at either speed (available only on the Black Diamond 8800 series switch, the BlackDiamond 12800 series switches, SummitStack, and the Summit family of switches)

In certain interoperability situations, it is necessary to turn autonegotiation off on a fiber gigabit Ethernet port. Even though a gigabit Ethernet port runs only at full duplex and gigabit speeds, the command that turns off autonegotiation must still include the duplex setting.

Gigabit Ethernet ports support flow control only when autonegotiation is turned on. When autonegotiation is turned off, flow control is not supported. (See the *ExtremeXOS Concepts Guide* for more detailed information on flow control on Extreme Networks devices.)

Example

The following example turns autonegotiation off for slot 2, port 1 on a modular switch:

```
configure ports 2:1 auto off speed 100 duplex full
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ports auto on

```
configure ports <port_list> auto on {speed [10 | 100 | 1000 | 10000]}
{duplex [full | half]}
```

Description

Enables autonegotiation for the particular port type.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
speed	Use this parameter to advertise the port speed as either 10, 100, 1000, or 10000 Mbps. .
duplex	Use this parameter to advertise the port duplex rate as either half or full.

Default

Auto on for 1 Gbps ports.

Auto off for 10 Gbps ports.

Usage Guidelines

The type of ports enabled for autonegotiation are 802.3u for 10/100 Mbps ports or 802.3z for gigabit Ethernet ports.

Flow control on gigabit Ethernet ports is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. (See the *ExtremeXOS Concepts Guide* for more detailed information on flow control on Extreme Networks devices.)

Example

The following command configures the switch to autonegotiate for slot 1, ports 2 and 4 on a modular switch:

```
configure ports 1:2, 1:4 auto on
```

History

This command was first available in ExtremeXOS 10.1.

The `speed` and `duplex` parameters were added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure ports auto-polarity

```
configure ports [<port_list> | all] auto-polarity [off | on]
```

Description

Configures the autopolarity detection feature on the specified Ethernet ports.

Syntax Description

port_list	Specifies one or more ports on the switch.
all	Specifies all of the ports on the switch.
off	Disables the autopolarity detection feature on the specified ports.
on	Enables the autopolarity detection feature on the specified ports.

Default

Enabled.

Usage Guidelines

This feature applies to only the 10/100/1000 BASE-T ports on the switch.

Use the `all` keyword to enable or disable the autopolarity detection feature on all of the Ethernet ports on the Summit family of switches, SummitStack, and the BlackDiamond 8800 series switch.

When autopolarity is disabled on one or more Ethernet ports, you can verify that status by using the command:

```
show ports information detail
```



NOTE

Autopolarity detection is always on for the BlackDiamond 10808 and BlackDiamond 12800 series switches.

Example

The following command disables the autopolarity detection feature on ports 5 to 7 on a Summit series switch:

```
configure ports 5-7 auto-polarity off
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the Summit family of switches, SummitStack, and the BlackDiamond 8800 series switch.

configure ports display-string

```
configure ports <port_list> display-string <string>
```

Description

Configures a user-defined string for a port or group of ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
string	Specifies a user-defined display string.

Default

N/A.

Usage Guidelines

The display string can be up to 15 characters. Display strings do not need to be unique for each port—you can assign the same string to multiple ports. For example, you could give all the ports that connected to a particular department a common display string.

The string is displayed in certain commands such as the `show ports information` command.



NOTE

Do not use a port number as a display string. For example, do not assign the display string "2" to port 2.

Example

The following command configures the user-defined string *corporate* for port 1 on a stand-alone switch:

```
configure ports 1 display-string corporate
```

The following command configures the user-defined string *corporate* for ports 3, 4, and 5 on slot 1 on a modular switch:

```
configure ports 1:3-5 display-string corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ports preferred-medium

```
configure ports <port_list> preferred-medium [copper | fiber] {force}
```

Description

Configures the primary uplink port to use a preferred medium.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
copper	Specifies that the port should always use the 10/100/1000 BASE-T connection whenever a link is established even when a fiber link is also present.
fiber	Specifies that the port should always use the 1 gigabit Ethernet fiber connection whenever a link is established even when a copper link is also present.
force	Disables automatic failover. (If the specified preferred medium is not present, the link does not come up even if the secondary medium is present.)

Default

The default is determined at cable installation time.

Usage Guidelines

You specify either `copper` or `fiber` for the specified port. The switch evaluates the copper energy and the fiber signal at the time these ports come online. If both are present, the configured preferred medium is chosen; however, if only one is present, the switch brings up that medium and uses this medium the next time the switch is rebooted. When a failure occurs and the uplinks are swapped, the switch continues to keep that uplink assignment until another failure occurs or until the assignment is changed using the CLI.

If you use the `force` option, it disables automatic failover. If you force the preferred-medium to fiber and the fiber link goes away, the copper link is not used, even if available.

To display the preferred medium, use the `show port information detail` command (you must use the `detail` variable to display the preferred medium).

Examples

The following establishes copper port 4 as the primary uplink on the Summit series switch and fiber port 4 as the redundant uplink port:

```
configure ports 4 preferred-medium copper
```

Copper port 4 becomes the primary uplink until a failure occurs on that link. At that time, fiber port 4 becomes the primary uplink and copper port 4 becomes the redundant port. This assignment stays in place until the next failure.

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the Summit family of switches, SummitStack, and the GM-20XT and GM-20XTR modules installed in a BlackDiamond 12800 series switch.

configure ports rate-limit flood

```
configure ports <port_list> rate-limit flood [broadcast | multicast |
unknown-destmac] [no-limit | <pps>]
```

Description

Limits the amount of ingress flooded traffic; minimizes network impact of broadcast loops.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
broadcast	Specifies all broadcast packets.
multicast	Specifies all flooded multicast packets (known IP multicast caches are still forwarded at line rate).
unknown-destmac	Specifies all packets with unknown MAC DAs.
no-limit	Specifies unlimited rate.
pps	Packets per second allowed; range is from 0 to 262,144.

Default

No limit.

Usage Guidelines

Use this command to limit the amount of ingress flooding traffic and to minimize the network impact of broadcast loops.

Example

The following command rate limits broadcast packets on port 3 on a stand-alone switch to 500 pps:

```
configure ports 3 rate-limit flood broadcast 500
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure ports redundant

```
configure ports <primaryPort> redundant <secondaryPort> {link [on | off]}
```

Description

Configures a software-controlled redundant port.

Syntax Description

primaryPort	Specifies one primary port or slot and port.
redundantPort	Specifies one or redundant port or slot and port.
link	Specifies state of link: - on—Specifies keeping the redundant port active, but block traffic - off—Specifies forcing the link down on the redundant port NOTE: The default value is off.

Default

N/A.

Usage Guidelines

The first port specifies the primary port. The second port specifies the redundant port.

A software-controlled redundant port is configured to back up a specified primary port; both ports are on the same device. The redundant port tracks the link state of the associated primary port, and if the link on the primary port fails, the redundant port establishes a link and becomes active. You can back up a specified Ethernet port with a redundant, dedicated Ethernet port.

You configure the redundant link to be always physically up but logically blocked or to be always physically down. The default is off, or the redundant link is down.

The following criteria must be considered when configuring a software-controlled redundant port:

- You can configure only one redundant port for each primary port.
- You cannot have any Layer 2 protocols configured on any of the VLANs that are present on the ports. (You will see an error message if you attempt to configure software redundant ports on ports with VLANs running Layer 2 protocols.)
- The primary and redundant port must have identical VLAN memberships.
- The master port is the only port of a load-sharing group that can be configured as either a primary or redundant port. (The entire trunk must go down before the software-controlled redundant port takes effect.)
- Only one side of the link should be configured as redundant.

**NOTE**

On the BlackDiamond 10808 switch, 10 Gbps modules with a serial number lower than 804405-00-09 the software redundant port feature cover only those failures where both the TX and RX paths fail. If a single strand of fiber is pulled on these ports, the software redundant port cannot correctly recover from the failure. To display the serial number of the module, issue the `show slot <slot_number>` command. (All the modules on the BlackDiamond 8800 series switch have this serial number or higher, as does the Summit family of switches.)

Example

The following command configures a software-controlled redundant port:

```
configure ports 1:3 redundant 2:3
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure ports wan-phy clocking

```
configure ports <port_list> wan-phy clocking [line | internal]
```

Description

Configures the clocking source for the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
line	Specifies the port clocking is recovered from the received bitstream.
internal	Specifies the port clock is an internal, free-run clock.

Default

Line.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

Each WAN PHY port can be configured for the source it will use for clocking; line clocking uses received clocking and internal clocking uses an internal clock.

Example

The following command configures an LW XENPAK WAN PHY port to use the free-run internal clock:

```
configure ports 1:3 wan-phy clocking internal
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on LW XENPAK ports on the BlackDiamond 10808 switch the Summit X450a series switches, and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

configure ports wan-phy framing

```
configure ports <port_list> wan-phy framing [sonet | sdh]
```

Description

Configures the framing type for the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
sonet	Specifies the port framing type as SONET.
sdh	Specifies the port framing type as SDH.

Default

SONET.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

Each WAN PHY port can be configured for framing that complies with either the SONET standard or the SDH standard. (SONET is primarily a U.S. standard, and SDH is the international version.)

Example

The following command configures an LW XENPAK WAN PHY port to use SONET framing:

```
configure ports 1:3 wan-phy framing sonet
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switches, and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

configure ports wan-phy loopback

```
configure ports <port_list> wan-phy loopback [line | off]
```

Description

Configures the loopback options for the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
line	Specifies the signal received from the optical media returns back to the transmitter.
off	Specifies no loopback.

Default

Off.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

Configuring loopback on a WAN PHY port is useful for diagnostics and network troubleshooting.

Example

The following command configures the loopback option on an LW XENPAK WAN PHY port as off:

```
configure ports 1:3 wan-phy loopback off
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switches, and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

configure ports wan-phy trace-path

```
configure ports <port_list> wan-phy trace-path <id_string>
```

Description

Configures the path trace identifier for the specified WAN PHY port. Path trace is a maintenance feature of WAN PHY.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
id_string	Enter an alphanumeric 16-character string.

Default

The IEEE default value, which has no string representation.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

The path trace message is used for troubleshooting. One byte of the path overhead associated with each WAN PHY interface SONET/SDH frame carries information identifying the originating path terminating equipment (PTE).

Example

The following command configures a patch trace ID for an LW XENPAK WAN PHY port:

```
configure ports 1:3 wan-phy trace-path bear3
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switches, and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

configure ports wan-phy trace-section

```
configure ports <port_list> wan-phy trace-section <id_string>
```

Description

Configures the section trace identifier for the specified WAN PHY port. Section trace is a maintenance feature of WAN PHY.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
id_string	Enter an alphanumeric 16-character string.

Default

The IEEE default value, which has no string representation.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

The section trace message is used for troubleshooting. The J0 transmit octets allow a receiver to verify its continued connection to the transmitter; the J0 octet transports a 16-octet continuously repeating section trace message.

The first transmitted section trace octet is J0 transmit 15, which contains the delineation octet; the default value is 137 (hexadecimal 89). The last transmitted section trace octet is J0 transmit 0; the default value for J0 transmit 0 through 14 is 0.

Example

The following command configures a section trace ID for an LW XENPAK WAN PHY port:

```
configure ports 1:3 wan-phy trace-section beta4
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switches, and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

configure sharing add ports

```
configure sharing <port> add ports <port_list>
```

Description

Adds ports to a load-sharing, or link aggregation, group. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the link aggregation group (LAG) if one port in the group goes down.

Syntax Description

port	Specifies the logical port for a load-sharing group or link aggregation group (LAG). This number also functions as the LAG Group ID.
port_list	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines



NOTE

You must create a LAG (or load-sharing group) before you can configure the LAG. To create a LAG, see [enable sharing <port> grouping <port_list> {algorithm \[port-based | address-based {L2 | L3 | L3_L4}\]}](#) {lACP}.

Use this command to dynamically add ports to a load-sharing group, or link aggregation group (LAG).

vMAN ports can belong to LAGs. If any port in the LAG is enabled for vMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, vMAN is automatically enabled on all ports of the untagged LAG.

To verify your configuration, use the [show ports sharing](#) command.

Summit family of switches only (does not apply to SummitStack). The following guidelines apply to link aggregation on the Summit family of switches:

- One static LAG can contain up to 8 ports.
- An LACP (dynamic) LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- The maximum number of LAGs is 128.
- The default load-sharing algorithm is L2 address-based aggregation. Any broadcast, multicast, or unknown unicast packet is transmitted on a single port in the LAG.

- The available address-based parameters on the Summit family of switches are as follows:
 - Summit X450 series switches—L2 for Layer 2 and L3 for Layer 3. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.
 - Summit X450a, X450e and X250e series switches—L2 for Layer 2, L3 for Layer 3, and L3_L4 for Layer 3/Layer 4.

**NOTE**

You cannot configure port-based load sharing algorithm on the Summit family of switches; you configure only address-based load-sharing algorithms.

BlackDiamond 8800 series switch and SummitStack only. The following guidelines apply to link aggregation on the BlackDiamond 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP (dynamic) LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- Any broadcast, multicast, or unknown unicast packet is transmitted on a single port in the LAG.

**NOTE**

You cannot configure port-based load sharing algorithm on the BlackDiamond 8800 series switch and the Summit family of switches; you configure only address-based load-sharing algorithms.

- The available address-based parameters on the BlackDiamond 8800 series switch are L2 for Layer 2 and L3 for Layer 3. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.
- Beginning with ExtremeXOS version 11.5 for the BlackDiamond 8800 series switch and ExtremeXOS version 12.0 for SummitStack, the maximum number of LAGs is 128 unless a 10G4X module is in use in a BlackDiamond 8800 series switch or a Summit X450-24t or X450-24x switch is in the SummitStack, in which case the maximum number of LAGs is 32.

**NOTE**

You cannot configure more than 32 LAGs on a BlackDiamond 8800 chassis with a 10G4X module running or that has a slot configured for the 10G4X module. You cannot configure more than 32 LAGs on a SummitStack that has a Summit X450-24t or X450-24x or that has a slot configured for either model.

If you attempt to configure more than 32 LAGs on a BlackDiamond 8800 chassis that contains an 10G4X module or has a slot configured for the 10G4X module, the system displays the following error message:

```
Error: Slot <slot_number> can support a maximum of 32 trunks
```

If you want to configure more than 32 LAGs, it is not necessary to remove the 10G4X or Summit X450 series module; however, you must *both* unconfigure the slot holding the module *and* disable the slot holding that module or remove the module (the disable slot command is not available on SummitStack). Use the following commands to unconfigure the slot and disable that slot:

```
unconfigure slot and disable slot
```


If you attempt to insert a 10G4X module into a BlackDiamond 8800 chassis or a SummitStack with a Summit X450 series module configured for more than 32 LAGs, the module fails the Init state with log error messages at the warning level similar to the following:

```
04/07/1921 23:52:28.29 <Warn:DM.Warning> MSM-A: Slot-8 FAILED (1) Error Max Load Share
Groups Exceeded(-48) from HAL on CardExec INIT(5) for slot
04/07/1921 23:52:28.29 <Warn:DM.Warning> MSM-A: Error Max Load Share Groups Exceeded(-
48) from HAL on CardExec INIT(5) for slot 8
```

Once you configure more than 32 LAGs, the module will not initialize even if you reduce the number of LAGs to 32; you must reboot the system or SummitStack first. The system logs an error message at the error level similar to the following when you must reduce the system (even if you reduced the number of LAGs to 32):

```
04/07/1921 23:52:28.29 <Erro:HAL.Card.Error> MSM-A: Slot 8 is not supported when more
than 32 load share groups have previously been configured. A system reboot is
required to clear this condition.
```

BlackDiamond 10808 and 12800 series switches only. The following guidelines apply to LAGs on the BlackDiamond 10808 and 12800 series switches:

- A static LAG can include a maximum 16 ports.
- An LACP (dynamic) LAG can include a maximum of 32 ports; out of these up to 16 can be selected links and the remaining 16 will be standby links.
- The maximum number of LAGs is 128.
- On the BlackDiamond 10808 and 12800 series switches, if you do not explicitly select an algorithm, the port-based scheme is used. However, the address-based algorithm has a more even distribution and is the recommended choice.



NOTE

You must use the `configure sharing address-based` command to configure address-based load sharing on the BlackDiamond 10808 and the BlackDiamond 12800 series switches.

Example

The following example adds port 3:13 to the LAG with the logical port 3:9 on a modular switch:

```
configure sharing 3:9 add port 3:13
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure sharing address-based

```
configure sharing address-based [L2 | L2_L3 | L2_L3_L4 | L2_L3_CHK_SUM |
L2_L3_L4_CHK_SUM]
```

Description

Configures the part of the packet examined by the switch when selecting the egress port for transmitting link aggregation, or load-sharing, data. Includes IPv6 addresses with ExtremeXOS version 11.2 and higher.

Syntax Description

L2	Indicates that the switch should examine the MAC source and destination address.
L2_L3	Indicates that the switch should examine the MAC source and destination address plus the IP source and destination address.
L2_L3_L4	Indicates that the switch should examine the MAC source and destination address, the IP source and destination, plus the TCP source and destination port number.
L2_L3_CHK_SUM	Indicates that the switch should examine the MAC source and destination address, the IP source and destination address, plus the IP check sum. NOTE: This variable applies only to IPv4 addresses.
L2_L3_L4_CHK_SUM	Indicates that the switch should examine the MAC source and destination address, the IP source and destination, the TCP source and destination port number, plus the IP check sum. NOTE: This variable applies only to IPv4 addresses.

Default

The default setting is L2.

Usage Guidelines



NOTE

You must create a LAG (or load-sharing group) before you can configure the LAG. To create a LAG, see [enable sharing <port> grouping <port_list> {algorithm \[port-based | address-based {L2 | L3 | L3_L4}}\] {lacp}](#).

This feature is available using the address-based load-sharing, or link aggregation, algorithm only. The address-based load-sharing algorithm uses addressing information to determine which physical port in the load-sharing group, or LAG, to use for forwarding traffic out of the switch. Addressing information is based on the packet protocol, as follows:

- IP packets—Uses the source and destination MAC and IP address, and the TCP port number.
- All other packets—Uses the source and destination MAC address.

This command specifies the part of the packet header that the switch examines to select the egress port for address-based load-sharing trunks. The address-based load-sharing setting is global and applies to all load-sharing trunks, or LAGs, that are address-based. You change this setting by issuing the command again with a different option.

To verify your configuration, use the `show ports sharing` command.

This is a global setting and applies to all load-sharing trunks that are address-based.

Because the algorithms `L2_L3__CHK_SUM` and `L2_L3_L4_CHK_SUM` use the IP check sum as part of the decision for the egress port selection, these algorithms do not guarantee the packet sequencing between the clients. Also, these variables apply only to IPv4 addresses.

If you do not explicitly select an algorithm, the port-based scheme is used. However, the address-based algorithm has a more even distribution and is the recommended choice.

Beginning with ExtremeXOS software version 11.3, you can configure dynamic link aggregation, using LACP.

Example

The following example configures the switch to examine the MAC source and destination address:

```
configure sharing address-based 12
```

History

This command was first available in ExtremeXOS 11.0.

IPv6-compatibility was added in ExtremeXOS 11.2.

The `L3_L4` optional parameter was added in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 10808 switch and the BlackDiamond 12800 series switches.

configure sharing delete ports

```
configure sharing <port> delete ports <port_list>
```

Description

Deletes ports from a link aggregation, or load-sharing, group.

Syntax Description

port	Specifies the logical port for a load-sharing group or a link aggregation group (LAG). This number also functions as the LAG Group ID.
port_list	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically delete ports from a load-sharing group, or link aggregation group (LAG). This command applies to static and dynamic link aggregation.

Example

The following example deletes port 3:12 from the LAG with the logical port, or LAG Group ID, 3:9:

```
configure sharing 3:9 delete port 3:12
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure sharing lacp activity-mode

```
configure sharing <port> lacp activity-mode [active | passive]
```

Description

Configures the whether the switch sends LACPDUs periodically (active) or only in response to LACPDUs sent from the partner on the link (passive).

Syntax Description

port	Specifies the master logical port for the LAG you are setting the activity mode for.
active	Enter this value to have the switch periodically sent LACPDUs for this LAG.
passive	Enter this value to have the switch only respond to LACPDUs for this LAG.

Default

Active.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP activity mode.



NOTE

One side of the link must be in active mode in order to pass traffic. If you configure your side in the passive mode, ensure that the partner link is in LACP active mode.

To verify the LACP activity mode, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different that a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



NOTE

In ExtremeXOS version 11.3, the activity mode cannot be changed from active.

Example

The following command changes the activity mode to passive for the specified LAG group ID:

```
configure sharing 5:1 lacp activity-mode passive
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure sharing lacp defaulted-state-action

```
configure sharing <port> lacp defaulted-state-action [add | delete]
```

Description

Configures whether a defaulted LAG port is removed from the aggregator.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the default action for.
add	Enter this value to have the switch add defaulted ports to the aggregator for this LAG.
delete	Enter this value to have the switch delete defaulted ports from the aggregator for this LAG.

Default

Delete.

Usage Guidelines

You must enable sharing and create the LAG prior to configuring this LACP parameter.

You can configure whether you want a defaulted LAG port removed from the aggregator or added back into the aggregator. If you configure the LAG to remove ports that move into the default state, those ports are removed from the aggregator and the port state is set to unselected.



NOTE

In ExtremeXOS version 11.3, defaulted ports in the LAG are always removed from the aggregator; this is not configurable.

If you configure the LAG to add the defaulted port into the aggregator, the system takes inventory of the number of ports currently in the aggregator:

- If there are fewer ports in the aggregator than the maximum number allowed, the system adds the defaulted port to the aggregator (port set to selected and collecting-distributing).
- If the aggregator has the maximum ports, the system adds the defaulted port to the standby list (port set to standby).



NOTE

If the defaulted port is assigned to standby, that port automatically has a lower priority than any other port in the LAG (including those already in standby).

To verify the LACP default action, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



NOTE

To force the LACP trunk to behave like a static sharing trunk, use this command to add ports to the aggregator.

Example

The following command deletes defaulted ports from the aggregator for the specified LAG group ID:

```
configure sharing 5:1 lacp defaulted-state-action delete
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure sharing lacp system-priority

```
configure sharing <port> lacp system-priority <priority>
```

Description

Configures the system priority used by LACP for each LAG to establish the station on which end assumes priority in determining those LAG ports moved to the collecting/distributing state of the protocol. That end of the LAG with the lowest system priority is the one that assumes control of the determination. This is optional; if you do not configure this parameter, LACP uses system MAC values to determine priority. If you choose to configure this parameter, enter a value between 1 and 65535.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the priority for.
priority	Enter the value you want for the priority of the system for the LACP. The range is 1 to 65535; there is no default.

Default

N/A.

Usage Guidelines

The LACP uses the system MAC values to assign priority to one of the systems, and that system then determines which LAG ports move into the collecting/distributing state and exchange traffic. That end of the LAG with the lowest system priority is the one that assumes control of the determination. If you wish to override the default LACP system priority for a specific LAG, use this command to assign that LAG a specific LACP priority. Enter a value between 1 and 65535.

You must enable sharing and create the LAG prior to assigning this LACP priority.

To verify the LACP system priority, use the `show lacp` command.

To change the system priority you previously assigned to a specific LAG, issue the `configure sharing lacp system-priority` using the new priority you want. To remove the assigned system priority entirely and use the LACP priorities, issue the `configure sharing lacp system-priority` using a value of 0.

Example

The following command assigns LAG 10 an LACP system priority of 3:

```
configure sharing 10 lacp system-priority 3
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure sharing lacp timeout

```
configure sharing <port> lacp timeout [long | short]
```

Description

Configures the timeout used by each LAG to stop transmitting once LACPDUs are no longer received from the partner link. You can configure this timeout value to be either 90 seconds, long, or 3 seconds, short.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the timeout value for.
long	Enter this value to use 90 seconds as the timeout value.
short	Enter this value to use 3 seconds as the timeout value.

Default

Long.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP timeout value.

To verify the LACP timeout value, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



NOTE

In ExtremeXOS version 11.3, the timeout value is set to long and cannot be changed.

Example

The following command changes the timeout value for the specified LAG group ID to short:

```
configure sharing 5:1 lacp timeout short
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure slot module

```
configure slot <slot> module <module_type>
```

Description

Configures a slot for a particular I/O module card in a modular switch.

On a stack, this command configures a slot for a particular type of node.

Syntax Description

slot	Specifies the slot number.
module_type	Specifies the type of module or node for which the slot should be configured. The list of modules you can enter will vary depending on the type of switch and version of ExtremeXOS you are running. Certain modules are supported only with specific ExtremeXOS Technology Releases. On a stack, module type will be the list of switches that support SummitStack.

Default

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Usage Guidelines

The `configure sharing lacp timeout` command displays different module parameters depending on the type of modular switch you are configuring and the version of ExtremeXOS running on the switch.

You can also preconfigure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated. If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type.

Upon powering up the chassis, or when an I/O module is hot-swapped, ExtremeXOS automatically determines the system power budget and protects the switch from any potential overpower configurations. If power is available, ExtremeXOS powers on and initializes the module. When ExtremeXOS detects that a module will cause an overpower condition, the module remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

On a stack, the module type must be a switch that supports SummitStack.

Example

The following command configures slot 2 for a 10/100/1000, 60-port, copper module:

```
configure slot 2 module G60T
```

The following command configures slot 2 for a X450a-24t on a stack:

```
configure slot 2 module X450a-24t
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

configure switch ports initial-mode disabled

```
configure switch ports <port_list> initial-mode disabled
```

Description

Disables default initialization mode of the switch.

Syntax Description

port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Ports are not added to the default VLAN on initial switch boot. Ports are disabled when the switch comes up. This command becomes active after the switch is rebooted with no configuration.

History

This command was available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 12800 series switches.

configure switch ports initial-mode enabled

```
configure switch ports <port_list> initial-mode enabled
```

Description

Enables the VLAN switch port initial configuration mode.

Syntax Description

port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Ports are not added to the default VLAN on initial switch boot. Ports are disabled when the switch comes up. This command becomes active after the switch is rebooted with no configuration.

History

This command was available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 12800 series switches.

disable edp ports

```
disable edp ports [<ports> | all]
```

Description

Disables the Extreme Discovery Protocol (EDP) on one or more ports.

Syntax Description

ports	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

You can use the `disable edp ports` command to disable EDP on one or more ports when you no longer need to locate neighbor Extreme Networks switches.

Example

The following command disables EDP on slot 1, ports 2 and 4 on a modular switch:

```
disable edp ports 1:2, 1:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable jumbo-frame ports

```
disable jumbo-frame ports [all | <port_list>]
```

Description

Disables jumbo frame support on a port.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports. NOTE: The BlackDiamond 8800 original series switch and the Summit X450 series switches (except the a- and e-series switches) support only all ports; you cannot enable or disable jumbo frames by port on the switch.

Default

Disabled.

Usage Guidelines

BlackDiamond 8800 original series switch modules and the Summit X450 series switches only. You can enable or disable jumbo frames for the entire module or switch on the following switches:

- BlackDiamond 8800 a-series and e-series modules
- Summit X450a, X450e and X250e series switches

These modules and switches do not support enabling or disabling jumbo frames by port. You must enable or disable jumbo frames for the entire switch.

BlackDiamond 8800 a- and e-series modules, BlackDiamond 10808 and BlackDiamond 12800 series switches, and Summit X250e, X450a, and X450e series switches only. You can enable jumbo frames by port.

Example

The following command disables jumbo frame support on slot 1, port 2 on a BlackDiamond 10808 switch:

```
disable jumbo-frame ports 1:2
```

The following command disables jumbo frame support on a BlackDiamond 8810 switch:

```
disable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable learning port

```
disable learning port [<port_list> | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.

Example

The following command disables MAC address learning on port 4:3 on a modular switch:

```
disable learning ports 4:3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable mirroring

```
disable mirroring
```

Description

Disables port mirroring.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use the `disable mirroring` command to stop all configured copied mirroring traffic. Use this command to unconfigure all the filters on the system.

Example

The following command disables port mirroring:

```
disable mirroring
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable port

```
disable port [<port_list> | all]
```

Description

Disables one or more ports on the switch.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command for security, administration, and troubleshooting purposes.

When a port is disabled, the link is brought down.

Example

The following command disables ports 3, 5, and 12 through 15 on a stand-alone switch:

```
disable ports 3,5,12-15
```

The following command disables slot 1, ports 3, 5, and 12 through 15 on a modular switch:

```
disable port 1:3,1:5,1:12-1:15
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable sharing

```
disable sharing <port>
```

Description

Disables a load-sharing group of ports, also known as a link aggregation group (LAG).

Syntax Description

port	Specifies the logical port of a load-sharing group or link aggregation group (LAG). Specifies a port or a combination of the slot and port number.

Default

Disabled.

Usage Guidelines

When sharing is disabled, the logical port retains all configuration including VLAN membership. All other member ports are removed from all VLANs to prevent loops and their configuration is reset to default values.

Example

The following command disables sharing on master logical port 9 in slot 3, which contains ports 9 through 12, on a modular switch:

```
disable sharing 3:9
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable slot

```
disable slot <slot> {offline}
```

Description

Disables slot and leaves that module in a power down state.

Syntax Description

slot	Specifies the slot to be disabled.
offline	Specifies that the slot be disabled offline.
NOTE: This variable is supported only on the BlackDiamond 8800 series switch; that is, those modular switches that support offline diagnostics.	

Default

Enabled.

Usage Guidelines

This command allows the user to disable a slot. When the user types this command, the I/O card in that particular slot number is brought down, and the slot is powered down. The LEDs on the card go OFF.

A disabled slot can be re-enabled using the `enable slot` command. When the slot is re-enabled, the software on the I/O module is updated to match the software on the primary MSM.

The `show slot` command, if invoked after the user disables the slot, shows this slot state as “Power Off/Disabled.”

If there is no I/O card present in a slot when the user disables the slot, the slot still goes to the “Disable” state. If a card is inserted in a slot that has been disabled, the card does not come up and stays in the “Power Off/Disabled” state until the slot is enabled by using the `enable slot` command. below.

If you do not save the configuration before you do a switch reboot, the slot will be re-enabled upon reboot. If you save the configuration after disabling a slot, the slot will remain disabled after a reboot.

On Power over Ethernet (PoE) modules, disabling a slot also disables any inline power that is flowing to that slot.

BlackDiamond 8800 series switch only. This command applies only to the data, or I/O ports on slots holding an MSM. The slots holding an MSM on the BlackDiamond 8810 switch are 5 and possibly 6; the slots holding an MSM on the BlackDiamond 8806 switch are 3 and possibly 4. Use the `offline` parameter to run the diagnostics offline.

Example

The following command disables slot 5 on the switch:

```
disable slot 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches. This command is not supported on a SummitStack.

disable smartredundancy

```
disable smartredundancy <port_list>
```

Description

Disables the Smart Redundancy feature.

Syntax Description

port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

The Smart Redundancy feature works in concert with the software-controlled redundant feature. When Smart Redundancy is disabled, the switch attempts only to reset the primary port to active if the redundant port fails. That is, if you disable Smart Redundancy, the traffic does not automatically return to the primary port once it becomes active again; the traffic continues to flow through the redundant port even after the primary port comes up again.

Example

The following command disables the Smart Redundancy feature on ports 1:1 to 1:4 on a modular switch:

```
disable smartredundancy 1:1-4
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable snmp traps port-up-down ports

```
disable snmp traps port-up-down ports [<port_list> | all]
```

Description

Disables port up/down trap reception for specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to stop receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command stops ports 3, 5, and 12 through 15 on a stand-alone switch from receiving SNMP trap messages when the port goes up/down:

```
disable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on all platforms.

enable edp ports

```
enable edp ports [<ports> | all]
```

Description

Enables the Extreme Discovery Protocol (EDP) on one or more ports.

Syntax Description

ports	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

On a modular switch or SummitStack, <ports> can be a list of slots and ports. On a stand-alone switch, <ports> can be one or more port numbers. For a detailed explanation of port specification, see [“Port Numbering”](#) in [Chapter 1, “Command Reference Overview.”](#)

EDP is useful when Extreme Networks switches are attached to a port.

The EDP is used to locate neighbor Extreme Networks switches and exchange information about switch configuration. When running on a normal switch port, EDP is used to by the switches to exchange topology information with each other. Information communicated using EDP includes the following:

- Switch MAC address (switch ID)
- Switch software version information
- Switch IP address
- Switch VLAN information
- Switch port number
- Switch port configuration data: duplex, and speed

Example

The following command enables EDP on slot 1, port 3 on a modular switch:

```
enable edp ports 1:3
```

History

This command was first available in ExtremeXOS 10.1.

The port configuration data was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable jumbo-frame ports

```
enable jumbo-frame ports [all | <port_list>]
```

Description

Enables support on the physical ports that will carry jumbo frames.

Syntax Description

all	Specifies ports.
port_list	Specifies one or more slots and ports. NOTE: The BlackDiamond 8800 original series switch and the Summit X450 and X250e series switch support only all ports; you cannot enable or disable jumbo frames by port on the switch.

Default

Disabled.

Usage Guidelines

Increases performance to back-end servers or allows for vMAN 802.1Q encapsulations.

You can configure the maximum size of a jumbo frame if you want to use a different size than the default value of 9216. Use the [configure jumbo-frame-size](#) command to configure the size.

This setting is preserved across reboots.

BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS 11.6, you can enable jumbo frames per port on the following switches:

- BlackDiamond a-series and e-series modules
- Summit X450a and X450e series switches and beginning with ExtremeXOS 12.0, the Summit X250e series switches

The original BlackDiamond 8800 series modules and the original Summit X450 series switches do not support enabling or disabling jumbo frames by port. If a SummitStack contains a Summit X450 series switch and Summit X450a, X450e, or X250e series switches, the entire SummitStack does not support enabling jumbo frames by port. You can enable or disable jumbo frames only for the entire switch. If you attempt to enable jumbo frames per port on a switch or SummitStack that does not support enabling jumbo frames by port, the system returns a platform-specific message similar to the following:

```
Error: You must enable jumbo-frame on all MSM-G8X I/O ports and G48T, G48P, G24x, and 10G4X ports globally. Use 'enable jumbo-frame ports all'.
```

If you attempt to enable jumbo frames per port on a BlackDiamond a-series or e-series module that co-exists in the chassis with an original BlackDiamond 8800 module, the system returns a platform-specific message similar to the following:

```
Error: You must first enable jumbo-frame on all MSM-G8X I/O ports and G48T, G48P, G24X, and 10G4X port globally. Use 'enable jumbo-frame ports all' and the 'disable jumbo-frame port <port_list>' on any other port.
```

After you issue the `enable jumbo-frames ports all` command, jumbo frames are enabled on all ports on any new modules that you bring online. You can then disable individual jumbo frames on individual ports on a- or e- series modules or switches.

BlackDiamond 10808 and the BlackDiamond 12800 series switches only. You can enable jumbo frames per port.

Example

The following command enables jumbo frame support on slot 3, port 5 on a BlackDiamond 10808 switch:

```
enable jumbo-frame ports 3:5
```

The following command enables jumbo frame support on a BlackDiamond 8810 switch:

```
enable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable learning port

```
enable learning port [all | <port_list>]
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

N/A.

Example

The following command enables MAC address learning on slot 1, ports 7 and 8 on a modular switch:

```
enable learning ports 1:7-8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable mirroring to port

```
enable mirroring to [port <port> [tagged | untagged] | port-list <port-list> loopback-port <port> [tagged | untagged]]
```

Description

Dedicates a port on the switch to be the mirror output port, or the monitor port.

Syntax Description

port	Specifies the mirror output port.
port-list	Specifies the list of ports where traffic is to be mirrored.
loopback-port	Specifies an otherwise unused port required when mirroring to a port-list. The loopback-port is not available for switching user data traffic.
port	Specifies a single loopback port that is used internally to provide this feature.
tagged	Specifies the mirror output, or monitor, port is tagged. NOTE: This parameter is available only on the BlackDiamond 10808 and the BlackDiamond 12800 series switches.
untagged	Specifies the mirror output, or monitor, port is untagged. NOTE: This parameter is available only on the BlackDiamond 10808 and the BlackDiamond 12800 series switches.

Default

Disabled.

Usage Guidelines

Port mirroring configures the switch to copy all traffic associated with one or more ports, VLANs or virtual ports. A virtual port is a combination of a VLAN and a port. The monitor port(s) can be connected to a network analyzer or RMON probe for packet analysis. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 16 mirroring filters and one monitor port can be configured on the switch. After a port has been specified as a monitor port, it cannot be used for any other function. Frames that contain errors are not mirrored.

You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

```
Error: Port mirroring cannot enabled on an ELSM enabled port 1.
```

Summit family of switches and SummitStack only. The traffic filter on the Summit family of switches can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port. You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.

- Egress—Mirrors traffic sent from the port.
- Ingress and egress—Mirrors traffic either received at the port or sent from the port.
(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).
- VLAN—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.
- Only 8 VLANs can be mirrored on a given physical port.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- In standard mode (see `configure mirroring mode` command), even if you select ingress and egress traffic, the packet is mirrored only the *first* time it matches a mirror filter and is not mirrored on subsequent configured filters. In enhanced mode, packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- You cannot include the monitor port for the Summit family of switches in a load-sharing group.
- You cannot run sFlow and mirroring on the same switch on a Summit X450 series switch. If you attempt to enable mirroring on these devices on a port that is already enabled for sFlow, the switch returns the following message:

```
Mirroring is not compatible with SFlow. Mirroring is not enabled!
```

You can run mirroring and sFlow on the same device when you are running one of the following:
 - Summit X450a, X450e or X250e series switch
- Tagged and untagged traffic is mirrored slightly differently depending on the switch that the mirrored port and the monitor port are on:
 - With a monitor port on a Summit X450 series switch (original Summit X450 switch), *all* traffic egressing the monitor port is tagged. Even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged with the internal VLAN ID.
 - With a monitor port on a Summit X450a, X450e or X250e series switch, the mirrored packet is tagged *only* if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port as untagged.
 - Egress mirrored packets are always transmitted as tagged on the monitor port(s).

BlackDiamond 8800 series switch only. The traffic filter on the BlackDiamond 8800 series switch can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port. You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors all traffic forwarded by the port.
(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).

- VLAN—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.
- Only 8 VLANs can be mirrored on a given physical port.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- In standard mode (see `configure mirroring mode` command), even if you select ingress and egress traffic, the packet is mirrored only the *first* time it matches a mirror filter and is not mirrored on subsequent configured filters. In enhanced mode, packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- You cannot include the monitor port for the BlackDiamond 8800 series switch in a load-sharing group.
- You cannot run sFlow and mirroring on the same switch for using the BlackDiamond 8800 original modules. If you attempt to enable mirroring on these devices on a port that is already enabled for sFlow, the switch returns the following message:

```
Mirroring is not compatible with SFlow. Mirroring is not enabled!
```

You can run mirroring and sFlow on the same device when you are running one of the following:
 - BlackDiamond a-series and e-series modules
- Tagged and untagged traffic is mirrored slightly differently depending on the module that the mirrored port and the monitor port are on:
 - With a monitor port on a BlackDiamond 8800 original module, *all* traffic egressing the monitor port is tagged (regardless of what module the ingressing port is on). Even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged with the internal VLAN ID.
 - With a monitor port on a BlackDiamond 8800 a-series or e-series module, the mirrored packet is tagged *only* if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port as untagged.
 - Egress mirrored packets are always transmitted as tagged on the monitor port(s).
- With the BlackDiamond 8800 a-series and e-series modules, you may see a packet mirrored twice. This occurs only if both the ingress mirrored port and the monitor port are on the same one-half of the module *and* the egress mirrored port is either on the other one-half of that module or on another module.

BlackDiamond 10808 and BlackDiamond 12800 series switches only. The traffic filter can be defined based on one of the following criteria:

- Physical port—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.
- VLAN—All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port—All data specific to a VLAN on a specific port is copied to the monitor port.

**NOTE**

Beginning with ExtremeXOS 11.0, the monitor port must be explicitly configured for tagged or untagged. Mirroring configurations prior to 11.0 are restored as tagged monitor ports only.

The traffic egressing the monitor port can either be tagged or untagged. If the mirroring is enabled as tagged on the monitor port of the BlackDiamond 10808 switch, all traffic egressing the monitor port is tagged. In this case, even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged. And, if mirroring is enabled as untagged on the monitor port, all traffic egressing the monitor port is untagged, including mirrored tagged packets.

When you upgrade to 11.0, all restored mirroring configurations are tagged monitor ports only.

Example

The following example selects slot 3, port 4 as the mirror, or monitor, port on the BlackDiamond 8810 switch:

```
enable mirroring to port 3:4
```

The following example selects slot 1, port 3 as the tagged mirror, or monitor, port on the BlackDiamond 10808 switch:

```
enable mirroring to port 1:3 tagged
```

History

This command was first available in ExtremeXOS 10.1.

Beginning with ExtremeXOS 11.0, you must configure the monitor port as tagged or untagged on the BlackDiamond 10808 and the BlackDiamond 12800 series switches.

Platform Availability

This command is available on all platforms.

enable port

```
enable port [<port_list> | all]
```

Description

Enables a port.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

All ports are enabled.

Usage Guidelines

Use this command to enable the port(s) if you disabled the port(s) for security, administration, or troubleshooting purposes.

Example

The following command enables ports 3, 5, and 12 through 15 on the stand-alone switch:

```
enable ports 3,5,12-15
```

The following command enables slot 1, ports 3, 5, and 12 through 15 on the modular switch:

```
enable port 1:3, 1:5, 1:12-1:15
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable sharing grouping

```
enable sharing <port> grouping <port_list> {algorithm [port-based |
address-based {L2 | L3 | L3_L4}]} {lacp}
```

Description

Enables the switch to configure port link aggregation, or load sharing. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the LAG if one port in the group goes down. LACP allows the system to dynamically configure the LAGs.

Syntax Description

port	Specifies the master logical port for a load-sharing group or link aggregation group (LAG).
port_list	Specifies one or more ports or slots and ports to be grouped to the logical port.
port-based	Specifies link aggregation by port-based algorithm. NOTE: This parameter is available only on the BlackDiamond 10808 and the BlackDiamond 12800 series switches.
address-based	Specifies link aggregation by address-based algorithm.
L2	Specifies address-based link aggregation by Layer 2. This is the default value. NOTE: This parameter is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches.
L3	Specifies address-based link aggregation by Layer 3. NOTE: This parameter is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit X450 series switch.
L3_L4	Specifies address-based link aggregation by Layer 3 IP plus Layer 4 port. NOTE: This parameter is available only on the BlackDiamond 8800 a-series and e-series switches and the Summit X450a, X450e and X250e series switches (whether or not included in a SummitStack).
lacp	Specifies dynamic link aggregation, or load sharing, using the LACP.

Default

Disabled.

Usage Guidelines

Link aggregation, or load sharing, allows you to increase bandwidth and availability between switches by using a group of ports to carry traffic in parallel between switches. The aggregation algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. Groups can span multiple modules.

Beginning with ExtremeXOS version 11.3, you can enable and configure dynamic link aggregation, using LACP. Static link aggregation is the default link aggregation method.

**NOTE**

Always verify the LACP configuration by issuing the `show ports sharing` command; look for the ports listed as being in the aggregator.

If a port in a LAG fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

Link aggregation must be enabled on both ends of the link, or a network loop will result.

**NOTE**

See ExtremeXOS Concepts Guide for information on the interaction of port-based ACLs and LAGs of ports.

Modular switch LAGs are defined according to the following rules:

- Although you can *reference* only the logical port of a LAG to a Spanning Tree Domain (STPD), *all* the ports of a load-sharing group actually belong to the specified STPD.
- When using link aggregation, you should always reference the logical port of the LAG when configuring or viewing VLANs. VLANs configured to use other ports in the LAG will have those ports deleted from the VLAN when link aggregation becomes enabled.

Link aggregation, or load-sharing, algorithms allow you to select the distribution technique used by the LAG to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering.

- Port-based—Uses the ingress port to determine which physical port in the LAG is used to forward traffic out of the switch.
- Address-based—Uses addressing information to determine which physical port in the LAG to use for forwarding traffic out of the switch. Refer to `configure sharing address-based` for more information on using addressing information.

Summit X450a X450e and X250e series switches and SummitStack only. The following guidelines apply to link aggregation on the Summit X450a, X450e and X250e series switches:

- One static LAG can contain up to 8 ports.
- An LACP (dynamic) LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- The maximum number of LAGs is 128.
- The default load-sharing algorithm is L2 address-based aggregation. Any broadcast, multicast, or unknown unicast packet is distributed across all members in the LAG.
- The available address-based parameters on the Summit X450a, X450e and X250e series switches are L2 for Layer 2 and L3_L4 for Layer 3 plus Layer 4. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.
- Beginning with ExtremeXOS version 12.0, the maximum number of LAGs in a SummitStack is 128 unless an original X450 series module is in use in the stack, in which case the maximum number of LAGs is 32.

**NOTE**

You cannot configure port-based load sharing algorithm on Summit X450 series switches; you configure only address-based load-sharing algorithms.

Summit X450 series switch only. The following guidelines apply to link aggregation on the Summit X450 series switch:

- One static LAG can contain up to 8 ports.
- An LACP (dynamic) LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- The maximum number of LAGs is 32.
- The default load-sharing algorithm is L2 address-based aggregation. Any broadcast, multicast, or unknown unicast packet is transmitted on a single port in the LAG.
- The available address-based parameters on the Summit X450 series switch are L2 for Layer 2 and L3 for Layer 3. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.

**NOTE**

You cannot configure port-based load sharing algorithm on the Summit X450 series switch; you configure only address-based load-sharing algorithms.

BlackDiamond 8800 series switch only. The following guidelines apply to link aggregation on the BlackDiamond 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP (dynamic) LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.

**NOTE**

You cannot configure port-based load sharing algorithm on the BlackDiamond 8800 series switch and the Summit X450 series switches; you configure only address-based load-sharing algorithms.

- The available address-based parameters on the BlackDiamond 8800 series switch are L2 for Layer 2 and L3 for Layer 3.
If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting. Beginning with ExtremeXOS software version 11.2, the switch can use IPv6 addresses.
- The maximum number of LAGs is 32.
- Broadcast, multicast, or unknown unicast packets are transmitted differently depending on the device you are using:
 - On the BlackDiamond 8800 original modules, these packets are transmitted on a single port of a LAG.
 - On the BlackDiamond 8800 a-series and e-series modules, these packets are distributed across all members of a LAG.
- Beginning with ExtremeXOS version 11.5, the maximum number of LAGs is 128 unless a 10G4X module is in use in the chassis, in which case the maximum number of LAGs is 32.

**NOTE**

You cannot configure more than 32 LAGs on a BlackDiamond 8800 chassis with a 10G4X module running or that has a slot configured for the 10G4X module.

The following switch configurations support 32 LAGs:

- A BlackDiamond 8800 switch that contains an 10G4X module or has a slot configured for the 10G4X module
- A SummitStack that contains an original Summit X450 switch or has a slot configured for the X450 switch

If you attempt to configure more than 32 LAGs on a switch or stack that does not support larger configurations, the system displays the following error message:

```
Error: Slot <slot_number> can support a maximum of 32 trunks
```

If you want to configure more than 32 LAGs on a BlackDiamond 8800 switch, it is not necessary to remove the 10G4X module; however, you must *both* unconfigure the slot holding the 10G4X module *and* disable the slot holding that module or remove the 10G4X module. Use the following commands to unconfigure the slot and disable that slot:

```
unconfigure slot and disable slot
```

If you want to configure more than 32 LAGs on a SummitStack, you must remove all original Summit X450 switches from the active topology.

If you attempt to insert a 10G4X module into a BlackDiamond 8800 chassis configured for more than 32 LAGs, or if you attempt to add an original X450 module into a SummitStack with more than 32 LAGs, the module fails the Init state with log error messages at the warning level similar to the following:

```
04/07/1921 23:52:28.29 <Warn:DM.Warning> MSM-A: Slot-8 FAILED (1) Error Max Load Share Groups Exceeded(-48) from HAL on CardExec INIT(5) for slot
04/07/1921 23:52:28.29 <Warn:DM.Warning> MSM-A: Error Max Load Share Groups Exceeded(-48) from HAL on CardExec INIT(5) for slot 8
```

Once you configure more than 32 LAGs on a BlackDiamond 8800 chassis or SummitStack, the 10G4X or original Summit X450 series module will not initialize even if you reduce the number of LAGs to 32; you must reboot the system or stack first. The system logs an error message at the error level similar to the following when you must reduce the system (even if you reduced the number of LAGs to 32):

```
04/07/1921 23:52:28.29 <Erro:HAL.Card.Error> MSM-A: Slot 8 is not supported when more than 32 load share groups have previously been configured. A system reboot is required to clear this condition.
```

When you have original BlackDiamond 8800 modules and BlackDiamond 8800 a-series and e-series modules in the same chassis, or when you have original X450 switches and X450a, X450e, or X250e switches in the same SummitStack, the configured algorithm may behave differently, as shown in the following:

Configured algorithm:	Original BlackDiamond 8800 module or Summit X450 series switch	BlackDiamond 8800 a-series and e-series modules or SummitStack with Summit X450a, X450e, or X250e switch
L2	L2	L2 (with flooding and multicast traffic balanced)
L3	L3	L3 and L4 ports
L3_L4	L3	L3 and L4 ports

BlackDiamond 10808 and BlackDiamond 12800 series switches only. The following guidelines apply to load sharing on the BlackDiamond 10808 and 12800 series switches:

- A static LAG can include a maximum 16 ports.
- An LACP (dynamic) LAG can include a maximum of 32 ports; out of these up to 16 can be selected links and the remaining 16 will be standby links.
- The maximum number of LAGs is 128.

On the BlackDiamond 10808 and 12800 series switches, if you do not explicitly select an algorithm, the port-based scheme is used. However, the address-based algorithm has a more even distribution and is the recommended choice.



NOTE

You must use the `configure sharing address-based` command to configure address-based load sharing on the BlackDiamond 10808 and 12800 series switches.

Example

The following example defines a static link aggregation group (LAG) on a modular switch that contains ports 9 through 12 on slot 3, ports 7 through 10 on slot 5, and uses the first port on slot 3 as the logical port 9 on a modular switch:

```
enable sharing 3:9 grouping 3:9-3:12, 5:7-5:10
```

In this example, logical port 3:9 represents physical ports 3:9 through 3:12 and 5:7 through 5:10.

The following example defines a dynamic LAG on a stand-alone switch containing ports 10 through 15, with port 10 being the logical port:

```
enable sharing 10 grouping 10-15 lacp
```

History

This command was first available in ExtremeXOS 10.1.

The address-based algorithm was added in ExtremeXOS 11.0.

The L2 and L3 optional parameters were added in ExtremeXOS 11.1.

IPv6-compatibility was added in ExtremeXOS 11.2.

Dynamic link aggregation, using LACP, was added in ExtremeXOS 11.3.

The L3_L4 optional parameter was added in ExtremeXOS 11.5.

SummitStack functionality was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable slot

```
enable slot <slot>
```

Description

Enables slots.

Syntax Description

slot	Specifies the slot to be enabled.

Default

Enabled.

Usage Guidelines



NOTE

On the BlackDiamond 8800 series switch, this command only applies to the data, or I/O, ports on slots holding an MSM (slot 5 and possibly 6 on the BlackDiamond 8810; slot 3 and possibly 4 on the BlackDiamond 8806 switch).

This command allows the user to enable a slot that has been previously disabled using the [disable slot](#) command.



NOTE

If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be enabled; the slot will not function in data-only mode without enough power for inline power.

When the user enters the enable command, the disabled I/O card in the specified slot is brought up, and the slot is made operational, if possible, or goes to the appropriate state as determined by the card state machine. The LEDs on the card are brought ON as usual. When the slot is enabled, the software on the I/O module is updated to match the software on the primary MSM.

After the user enables the slot, the [show slot](#) command shows the state as “Operational” or will display the appropriate state if the card could not be brought up successfully. Note that there is no card state named “Enable” and the card goes to the appropriate states as determined by the card state machine when the [enable slot](#) command is invoked.

Only slots that have their state as “disabled” can be enabled using this command. If this command is used on slots that are in states other than “disabled,” the card state machine takes no action on these slots.

To enable inline power to a slot, the slot must be enabled as well as inline power for that slot. Use the [enable inline-power](#) command to enable inline power.

Example

The following command enables slot 5 on the switch:

```
enable slot 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

enable smartredundancy

```
enable smartredundancy <port_list>
```

Description

Enables the Smart Redundancy feature on the primary port.

Syntax Description

portlist	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

You must configure the software-controlled redundant port using the `configure ports redundant` command prior to enabling Smart Redundancy.

The Smart Redundancy feature works in concert with the software-controlled redundant port feature. With Smart Redundancy enabled on the switch, when the primary port becomes active the switch redirects all traffic to the primary port and blocks the redundant port again. (If you disable Smart Redundancy, the primary port is blocked because traffic is now flowing through the redundant, port.)

Example

The following command enables the Smart Redundancy feature on slot 1, port 4 on a modular switch:

```
enable smartredundancy 1:4
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable snmp traps port-up-down ports

```
enable snmp traps port-up-down ports [<port_list> | all]
```

Description

Enables port up/down trap reception for specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to begin receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command enables ports 3, 5, and 12 through 15 on a stand-alone switch to receive SNMP trap messages when the port goes up/down:

```
enable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on all platforms.

restart ports

```
restart ports [all | <port_list>
```

Description

Resets autonegotiation for one or more ports by resetting the physical link.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command resets autonegotiation on slot 1, port 4 on a modular switch:

```
restart ports 1:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

run failover

```
run failover {force}
```

Description

Causes a user-specified node failover.

Syntax Description

force	Force failover to occur.

Default

N/A.

Usage Guidelines

Use this command to cause the primary MSM to failover to the backup MSM on modular switches, or the Master node to failover to the Backup node in SummitStack.

Before you initiate failover, use the `show switch {detail}` command to confirm that the nodes are in sync and have identical software and switch configurations. If the output shows MASTER and BACKUP (InSync), the two MSMs or nodes are in sync.

If the MSMs software and configuration are not in sync, and both MSMs are running ExtremeXOS 11.0 or later, or the master and backup SummitStack nodes software and configuration are not in sync and are running ExtremeXOS 12.0 or later, use the `synchronize` command to get the two MSMs or nodes in sync. This command ensures that the backup has the same software in flash as the master.



NOTE

Both the backup and the master MSMs or nodes must be running ExtremeXOS 11.0 or later to use the `synchronize` command.

If the MSMs are not in sync, and one MSM is running ExtremeXOS 10.1 or earlier, specify the `force` option of the `run failover` command. By specifying `force`, failover occurs regardless of the version of software running on the MSMs.

Example

The following command causes a failover on a modular switch or SummitStack:

```
run failover
```

History

This command was first available for modular switches and SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

run msm-failover

```
run msm-failover {force}
```

Description

Causes a user-specified node failover.

This command is not supported on a SummitStack. To do a failover on a stack, use the command `run failover`.

Syntax Description

force	Force failover to occur.

Default

N/A.

Usage Guidelines

This command is being replaced with the `run failover` command. For usage guidelines, see the description for the `run failover` command.

Example

The following command causes a user-specified MSM failover:

```
run msm-failover
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

show edp

```
show edp {ports [all | <ports>] {detail}}
```

Description

Displays connectivity and configuration information for neighboring Extreme Networks switches.

Syntax Description

ports	Specifies one or more ports or slots and ports.
all	Specifies all ports.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

On a modular switch, <ports> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see [“Port Numbering”](#) in [Chapter 1, “Command Reference Overview.”](#) To clear the counters, use the `clear lacp counters` command.

The neighbor-ID value is eight bytes. The first two bytes are always set to 00:00; the last six bytes are set to the neighbor’s system MAC address.

Use the `show edp` command to display neighboring switches and configurations. This is most effective with Extreme Networks switches.

Example

The following command displays the configuration of the switch:

```
show edp
```

Following is sample output from this command:

```
EDP advert-interval      :60 seconds
EDP holddown-interval    :180 seconds
EDP enabled on ports     :1:1  1:2  1:3  1:4  1:5  1:6  3:1  3:2  3:3  3:4
```

Following is sample output from the `show edp ports 1:1` command:

```
Port   Neighbor                Neighbor-ID                Remote  Age   Num
      Port                  Port                      Port    Vans
=====
1:1    Oban                    00:00:00:30:48:41:ed:97   1:1    54    1
=====
```

The following command displays the connectivity and configuration of neighboring Extreme Networks switches:

```
show edp ports 1:1 detail
```

Following is sample output from this command:

```
=====
Port 1:1: EDP is Enabled
Tx stats: sw-pdu-tx=2555      vlan-pdu-tx=1465      pdu-tx-err=0
Rx stats: sw-pdu-rx=2511      vlan-pdu-rx=2511      pdu-rx-err=0

Time of last transmit error: None
Time of last receive error:  None
Remote-System:                Oban                        Age = 41
Remote-ID:                    00:00:00:30:48:41:ed:97
Software version:             11.1.0.19
Remote-Port:                  1:1
Port Type:                    Ethernet
Auto Negotiation:             OFF
Flow Control:                 SYMMETRIC/ASYMMETRIC
Duplex Speed:                 Configured = HALF          Actual = HALF
Port Speed (MB):              Configured = ERROR         Actual = 100 Mbps
Remote-Vlans:
    test (4094)  Age = 41
=====
```



NOTE

The output differs if the port is connected to a port running ExtremeWare software; the output shown above is displayed when both connected ports are running ExtremeXOS software.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show lacp

```
show lacp
```

Description

Displays LACP, or dynamic link aggregation, settings on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information about the LACP LAGs configured on the switch:

- Up or Down
- Enabled or disabled (not configurable)
- System MAC
 - MAC address for the system, which is used for LACP priority in the absence of a specifically configured priority.
- LACP PDUs dropped on non-LACP ports
- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.

- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.

Example

The following command displays the LACP LAGs on the switch:

```
show lacp
```

The following is sample output from this command on a modular switch:

```
LACP Up : Yes
LACP Enabled : Yes
System MAC : 00:04:96:10:33:60
LACP PDUs dropped on non-LACP ports : 0
```

Lag	Actor Sys-Pri	Actor Key	Partner MAC	Partner Sys-Pri	Partner Key	Agg Count
2:1	90	0x07d1	00:01:30:f9:9c:30	601	0x1391	2
4:5	100	0x0fa5	00:01:30:f9:9c:30	321	0x1f47	16
4:9	677	0x0fa9	00:01:30:f9:9c:30	87	0x0fa9	8

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show lacp counters

```
show lacp counters
```

Description

Displays all LACP, or dynamic link aggregation, counters for all member ports in the system.

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command displays the following information for all link aggregation groups (LAGs):

- LACP PDUs dropped on non-LACP ports
- LACP bulk checkpointed messages sent
- LACP bulk checkpointed messages received
- LACP PDUs checkpointed sent
- LACP PDUs checkpointed received
- LAG group ID
- Member port
- Packets received
- Packets dropped from PDU error
- Packets dropped because LACP is not enabled on this port
- Packets dropped because sender's system MAC address matches that of receiver
- Packets successfully transmitted
- Packets with errors during transmission

Example

The following command displays LACP counters:

```
show lacp counters
```

The following is sample output from this command on a modular switch:

```
LACP PDUs dropped on non-LACP ports : 519392
LACP Bulk checkpointed msgs sent    : 1
LACP Bulk checkpointed msgs recv   : 0
LACP PDUs checkpointed sent        : 575616
LACP PDUs checkpointed recv        : 0
```

Lag Group	Member Port	Rx Ok	Rx Drop PDU	Rx Drop Err	Rx Drop Not Up	Rx Drop Same MAC	Tx Sent	Tx Ok	Tx Xmit	Tx Err
1:1	1:1	2169	0	0	0	0	2170	0		
	1:2	2169	0	0	0	0	2170	0		
	1:3	2169	0	0	0	0	2170	0		
	1:4	2169	0	0	0	0	2170	0		
	1:5	2169	0	0	0	0	2170	0		
	1:6	2169	0	0	0	0	2170	0		
	1:7	2169	0	0	0	0	2170	0		
	1:8	2168	0	0	0	0	2169	0		
=====										

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show lacp lag

```
show lacp lag <group-id> {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified LAG.

Syntax Description

group-id	Specifies the LAG group ID you want to display. This is the number of the port you configured as the logical port of the LAG.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified LACP LAG:

- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.
- Member port

- Port priority
- Rx State—Receiving state of the port
 - Idle
 - Initialized
 - Current—Receiving LACP PDUs
 - Expired
 - Defaulted
- Sel Logic—Selection state of the port
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag—Mux state of the port
 - A—Activity
 - T—Timeout
 - G—Aggregation
 - S—Synchronization
 - C—Collecting
 - D—Distributing
 - F—Defaulted
 - E—Expired
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up—Yes or no
- Enabled—Yes or no
- Unack count
- Wait-for-count
- Current timeout
- Activity mode
- Defaulted action
- Receive state
- Transmit state
- Selected count—Number of selected ports in the LAG

- Standby count—Number of standby ports in the LAG
- LAG Id flag
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.

Example

The following command displays information on the specified LACP LAG:

```
show lacp lag 4:9
```

The following is sample output from this command on a modular switch:

Lag	Actor	Actor	Partner	Partner	Partner	Agg
Sys-Pri	Key	MAC	Sys-Pri	Key	Count	
4:9	2110	0x0fa9	00:04:96:10:33:60	2110	0x0fa9	16

Port list:

Member	Port	Rx	Sel	Mux	Actor	Partner
Port	Priority	State	Logic	State	Flags	Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009
4:10	301	Current	Selected	Collect-Dist	A-GSCD--	4010
4:11	302	Current	Standby	Detached	A-G-----	4011
4:12	303	Current	Standby	Detached	A-G-----	4012
4:29	200	Current	Selected	Collect-Dist	A-GSCD--	4029
4:30	0	Current	Selected	Collect-Dist	A-GSCD--	4030
4:31	202	Current	Selected	Collect-Dist	A-GSCD--	4031
4:32	203	Current	Selected	Collect-Dist	A-GSCD--	4032
8:7	101	Current	Selected	Collect-Dist	A-GSCD--	8013
8:8	10	Current	Selected	Collect-Dist	A-GSCD--	8014
8:9	9	Current	Selected	Collect-Dist	A-GSCD--	8015
8:10	8	Current	Selected	Collect-Dist	A-GSCD--	8016
8:11	7	Current	Selected	Collect-Dist	A-GSCD--	8017
8:12	6	Current	Selected	Collect-Dist	A-GSCD--	8018
8:13	5	Current	Selected	Collect-Dist	A-GSCD--	8019
8:14	3	Current	Selected	Collect-Dist	A-GSCD--	8020
8:15	0	Current	Selected	Collect-Dist	A-GSCD--	8043
8:16	3	Current	Selected	Collect-Dist	A-GSCD--	8044
8:17	2	Idle	Unselected	Detached	-----	0
8:18	37	Idle	Unselected	Detached	-----	0
8:19	36	Idle	Unselected	Detached	-----	0
8:20	35	Idle	Unselected	Detached	-----	0

=====

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

The following command displays detailed information on the specified LACP LAG:

```
show lacp lag 4:9 detail
```

The following is sample output from this command on a modular switch:

```

Lag          Actor      Actor      Partner      Partner  Partner  Agg
            Sys-Pri   Key        MAC          Sys-Pri   Key      Count
-----
4:9          2110      0x0fa9     00:04:96:10:33:60  2110      0x0fa9   16

Up           : Yes
Enabled      : Yes
Unack count  : 0
Wait-for-count : 0
Current timeout : Long
Activity mode : Active
Defaulted Action : Delete
Receive state : Enabled
Transmit state : Enabled
Selected count : 16
Standby count : 2
LAG Id flag   : Yes
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9

```

Port list:

```

Member      Port      Rx      Sel      Mux      Actor      Partner
Port        Priority  State   Logic     State     Flags      Port
-----
4:9          300      Current Selected Collect-Dist A-GSCD-- 4009
4:10         301      Current Selected Collect-Dist A-GSCD-- 4010
4:11         302      Current Standby   Detached  A-G----- 4011
4:12         303      Current Standby   Detached  A-G----- 4012
4:29         200      Current Selected Collect-Dist A-GSCD-- 4029
4:30          0      Current Selected Collect-Dist A-GSCD-- 4030
4:31         202      Current Selected Collect-Dist A-GSCD-- 4031
4:32         203      Current Selected Collect-Dist A-GSCD-- 4032
8:7          101      Current Selected Collect-Dist A-GSCD-- 8013
8:8           10      Current Selected Collect-Dist A-GSCD-- 8014
8:9           9      Current Selected Collect-Dist A-GSCD-- 8015
8:10          8      Current Selected Collect-Dist A-GSCD-- 8016
8:11          7      Current Selected Collect-Dist A-GSCD-- 8017
8:12          6      Current Selected Collect-Dist A-GSCD-- 8018
8:13          5      Current Selected Collect-Dist A-GSCD-- 8019
8:14          3      Current Selected Collect-Dist A-GSCD-- 8020
8:15          0      Current Selected Collect-Dist A-GSCD-- 8043
8:16          3      Current Selected Collect-Dist A-GSCD-- 8044
8:17          2      Idle     Unselected Detached  ----- 0
8:18         37      Idle     Unselected Detached  ----- 0
8:19         36      Idle     Unselected Detached  ----- 0
8:20         35      Idle     Unselected Detached  ----- 0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show lacp member-port

```
show lacp member-port <port> {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified port that is a member of any LAG.

Syntax Description

port	Specifies the port number.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified port:

- Member Port
- Port Priority
- Rx State—Receiving state of the port
 - Idle
 - Initialized
 - Current—Receiving LACP PDUs
 - Expired
 - Defaulted
- Sel Logic—Selection state of the port
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.

- Actor Flag
 - A—Activity
 - T—Timeout
 - G—Aggregation
 - S—Synchronization
 - C—Collecting
 - D—Distributing
 - F—Defaulted
 - E—Expired
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up or Down—LACP protocol running or not on specified port
- Enabled or disabled (not configurable)
- Link State—Link state on this port up or down
- Actor Churn—True or false
- Partner Churn—True or false
- Ready_N—Ready to be added to aggregator.
- Wait pending
- Ack pending
- LAG Id
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.
- Stats
 - Rx - Accepted
 - Rx - Dropped due to error in verifying PDU
 - Rx - Dropped due to LACP not being up on this port
 - Rx - Dropped due to matching own MAC
 - Tx - Sent Successfully
 - Tx - Transmit error

Example

The following command displays LACP information on the specified port:

```
show lacp member-port 4:9
```

The following is sample output from this command on a modular switch:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

=====

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

The following command displays detailed LACP information on the specified port:

```
show lacp member-port 4:9 detail
```

The following is sample output from this command on a modular switch:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

```

Up : Yes
Enabled : Yes
Link State : Up
Actor Churn : False
Partner Churn : False
Ready_N : Yes
Wait pending : No
Ack pending : No
LAG Id:
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9, P.pri:300 , P.num:4009
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9, Q.pri:300 , Q.num:4009
Stats:
Rx - Accepted : 2174
Rx - Dropped due to error in verifying PDU : 0
Rx - Dropped due to LACP not being up on this port : 0
Rx - Dropped due to matching own MAC : 0

Tx - Sent successfully : 2175
Tx - Transmit error : 0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
             C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show mirroring

```
show mirroring
```

Description

Displays the port-mirroring configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must enable mirroring on the switch prior to configuring mirroring, and you must configure mirroring to display mirroring statistics. Use the `enable mirroring to port` command to enable mirroring and the `configure lacp member-port priority` command to configure mirroring.

You can use this command to display mirroring statistics and determine if mirroring is enabled or disabled on the switch.

Example

The following command displays switch mirroring statistics:

```
show mirroring
```

BlackDiamond 8800 series switch and Summit family of switches only.

Following is sample output from this command for a BlackDiamond 8810 switch or SummitStack that is configured for port-based mirroring for single monitor ports:

```
Mirror port: 3:15 is up
Number of Mirroring filters: 3
Mirror Port configuration:
    Port number 3:12 in  all vlans ingress only
    Port number 5:4 in  all vlans egress only
    Port number 8:30 in  all vlans
```

Following is sample output from this command on the Summit X450 series switch that is configured for mirroring a virtual port:

```
Mirror port: 12 is down
Number of Mirroring filters: 1
Mirror Port configuration:
    Port number 3 in vlan peggy.
```


Following is a sample output from this command for a BlackDiamond 8800 switch, SummitStack or Summit X450 switch as enabled for multiple monitor ports.

```
Mirror ports: 5-7 Loopback port: 3
Number of Mirroring filters: 2
Mirror Port configuration:
    Port number 1 in all vlans
    Port number 2 in all vlans
```

BlackDiamond 10808 and BlackDiamond 12800 series switches only.

Following is sample output from this command on the BlackDiamond 10808 and 12800 series switches that is configured for port-based mirroring for single monitor ports:

```
Mirror port: 1:5 is up tagged
Number of Mirroring filters: 1
Mirror Port configuration:
    Port number 2:1 in all vlans
```

Following is sample output from this command on the BlackDiamond 10808 and 12800 series switches as enabled for multiple monitor ports.

```
Mirror ports: 1:2-16 tagged Loopback port: 1:1
Number of Mirroring filters: 1
Mirror Port configuration:
    Port number 1:40 in all vlans
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ports anomaly

```
show ports <port list> anomaly {no-refresh}
```

Description

Display statistics of anomaly violation events in real time.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

N/A.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This command takes effect after enabling anomaly-protection.

Example

The following command displays real-time anomaly statistics on slot 2, all ports on a modular switch:

```
show ports 2:* anomaly
```

Following is sample output from this command:

```
Port Statistics Thu Nov  9 22:44:31 2006
Port   Link      Rx Pkt ===== Anomaly Violation =====
      Status    Count          L3 Count          L4 Count          ICMP Count          Frag Count
=====
2:1     A        191585             1              2              0              0
2:2     R           0              0              0              0              0
2:3     R           0              0              0              0              0
2:4     R           0              0              0              0              0
2:5     R           0              0              0              0              0
2:6     R           0              0              0              0              0
2:7     R           0              0              0              0              0
2:8     R           0              0              0              0              0
2:9     R           0              0              0              0              0
2:10    R           0              0              0              0              0
2:11    R           0              0              0              0              0
2:12    A       178024             0              0              0              0
```

2:13	A	196956	0	0	0	0
2:14	R	0	0	0	0	0
2:15	R	0	0	0	0	0
2:16	R	0	0	0	0	0
2:17	R	0	0	0	0	0

=====

Link Status : A-Active, R-Ready, NP-Port Not Present L-Loopback
0->Clear Counters U->page up D->page down ESC->exit

show ports collisions

```
show ports {mgmt | <port_list>} collisions {no-refresh}
```

Description

Displays real-time collision statistics.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, collision statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays real-time collision statistics on slot 1, ports 1 and 2 on a modular switch:

```
show ports 1:1-2 collisions
```

Following is sample output from this command:

```
Port Collision Monitor
Port      Link      Collision Histogram
Status 1   2   3   4   5   6   7   8   9  10  11  12  13  14  15  16
=====
1:1      A      0   0   0   0   0   0   0   0   0   0   0   0   0   0   0
1:2      R      0   0   0   0   0   0   0   0   0   0   0   0   0   0   0
=====
Link Status: A-Active R-Ready, NP-Port not present, L-Loopback
```

The numbers 1 to 16 represent the number of collisions encountered prior to successfully transmitting the packet; this is applicable only for half-duplex links.

History

This command was first available in ExtremeXOS 10.1.

The `Not Present` variable was added in ExtremeXOS 11.2.

The `no-refresh` variable was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show ports configuration

```
show ports {mgmt | <port_list>} configuration {no-refresh}
```

Description

Displays port configuration statistics, in real time or snapshot.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, configuration statistics are displayed for all ports. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This status information may be useful for your technical support representative if you have a network problem.

This command displays port configuration, which includes:

- Virtual router
- Port state
- Autonegotiation information
- Link state
- Link speed
- Duplex mode
- Flow control
- Load sharing information
- Link media information



NOTE

On 10 Gbps ports, the Media Primary column displays *NONE* when no module is installed, and *SR*, *LR*, or *ER* depending on the module installed when there is one present.

Example

The following command displays the port configuration for all ports on a modular switch:

```
show ports configuration
Port Configuration Monitor                               Fri Apr 13 10:22:29 2007
Port      Virtual      Port Link Auto   Speed      Duplex  Flow  Load  Media
          router        State State Neg   Cfg Actual Cfg Actual Cntrl Master Pri Red
=====
1         VR-Default    E     R    ON   AUTO      AUTO                      NONE UTP
2         VR-Default    E     R    ON   AUTO      AUTO                      NONE UTP
3         VR-Default    E     R    ON   AUTO      AUTO                      NONE UTP
4         VR-Default    E     R    ON   AUTO      AUTO                      NONE UTP
5         VR-Default    E     R    ON   AUTO      AUTO                      NONE
6         VR-Default    E     R    ON   AUTO      AUTO                      NONE
7         VR-Default    E     R   OFF   100      FULL                      SX
8         VR-Default    E     R    ON   AUTO      AUTO                      NONE
9         VR-Default    E     R    ON   AUTO      AUTO                      NONE
10        VR-Default    E     R    ON   AUTO      AUTO                      NONE
11        VR-Default    E     R    ON   AUTO      AUTO                      NONE
12        VR-Default    E     R    ON   AUTO      AUTO                      NONE
13        VR-Default    E     R    ON   AUTO      AUTO                      NONE
14        VR-Default    E     R    ON   AUTO      AUTO                      NONE
15        VR-Default    E     R    ON   AUTO      AUTO                      NONE
=====
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, Media: !-Unsupported Optic Module
Media Red: * - use "show port info detail" for redundant media type
0->Clear Counters  U->page up  D->page down  ESC->exit
```

The following command displays the port configuration statistics for slot 2, port 2:

```
show ports 2:2 configuration
```

Following is sample output from this command:

```
Port Configuration
Port      Virtual      Port Link Auto   Speed      Duplex  Flow  Load  Media
          router        State State Neg   Cfg Actual Cfg Actual Cntrl Master Pri Red
=====
2:2       VR-Default    E     R    ON   AUTO      AUTO                      UTP
=====
Link Status: A-Active, R-Ready, NP-Port not present, L-Loopback
Port State: D-Disabled E-Enabled, Media: !-Unsupported Optic Module
0->Clear Counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

The Port not present and Media variables were added in ExtremeXOS 11.2.

The no-refresh variable was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show ports information

```
show ports {mgmt | <port_list>} information {detail}
```

Description

Displays detailed system-related information.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports of slots and ports.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

This command displays information, including the following:

- Port number
- Port configuration
 - Virtual router
 - Type of port
 - Admin state
 - Link state and speed
 - Link counter
 - VLAN configuration
 - STP configuration
 - Trunking, or load sharing
 - EDP
 - ELSM (disabled; or if enabled, the ELSM link state is shown as well)
 - Load balancing
 - Learning
 - Egress flooding
 - Jumbo frames
 - Link port up/down traps
 - QoS profiles
 - vMAN status
 - Smart Redundancy status
 - SRP status

- Additional platform-specific information
- LW XENPAK WAN PHY ports—BlackDiamond 10808 switch and Summit X450a series switch only
 - Framing
 - Clocking
 - Trace section
 - Trace path

If you do not specify a port number or range of ports, detailed system-related information is displayed for all ports. The data is displayed in a table format.

This status information may be useful for your technical support representative if you have a network problem.

The `detail` parameter is used to provide more specific port information. The data is called out with written explanations versus displayed in a table format.



NOTE

The keyword `detail` displays slightly different information depending on the platform and configuration you are working with.

The link filter counter displayed with the `detail` keyword is calculated at the middle layer on receiving an event. The link filter up indicates the number of link transitions from down to up at the middle layer filter.

Example

The following command displays port system-related information on a BlackDiamond 8810 switch:

```
show port 3:1 information
```

Following is sample output from this command:

```
BD8810.4 > show port 3:1 info
```

```

Port          Flags          Link          Link Num Num  Num  Jumbo QOS      Load
              State ELSM UPS   STP VLAN Proto Size  profile Master
=====
3:1           Em-----e--  ready   -   0   0   1     1  9216
=====
Flags : a - Load Sharing Algorithm address-based, D - Port Disabled,
        e - Extreme Discovery Protocol Enabled, E - Port Enabled,
        g - Egress TOS Enabled, j - Jumbo Frame Enabled,
        l - Load Sharing Enabled, m - MACLearning Enabled,
        n - Ingress TOS Enabled, o - Dot1p Replacement Enabled,
        P - Software redundant port(Primary),
        R - Software redundant (Redundant),
        q - Background QOS Monitoring Enabled,
        s - diffserv Replacement Enabled,
        v - Vman Enabled, f - Unicast Flooding Enabled
        M - Multicast Flooding Enabled, B - Broadcast Flooding Enabled
        L - Extreme Link Status Monitoring Enabled

```

**NOTE**

The BlackDiamond 10808 series and 12800 series switches have an additional flag: p - Load Sharing Algorithm port-based.

BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches only. The following command displays more specific information for slot 3, port 1 on a BlackDiamond 8810 switch:

```
show ports 3:1 information detail
```

Following is sample output from this command:

```
Port:      3:1
Virtual-router: VR-Default
Type:      UTP
Random Early drop:      Unsupported
Admin state:      Enabled with auto-speed sensing (100M Advertised), auto-
duplex (half-duplex Advertised)
ELSM Link State:      Up
Link State:      Active, 1 Gbps, full-duplex
Link Counter: Up      1 time(s)
VLAN cfg:
    Name: Default, Internal Tag = 1 (MAC-Based), MAC-limit = No-limit

STP cfg:
    s0(disable), Tag=(none), Mode=802.1D, State=FORWARDING

Protocol:
    Name: Default      Protocol: ANY      Match all protocols.
Trunking:      Load sharing is not enabled.
EDP:      Enabled
ELSM:      Enabled
Learning:      Enabled
Unicast Flooding:      Enabled
Multicast Flooding:      Enabled
Broadcast Flooding:      Enabled
Jumbo:      Enabled, MTU= 9216
Link up/down SNMP trap filter setting:      Enabled
Egress Port Rate:      128 Kbps, Max Burst Size: 200 Kb
Broadcast Rate:      No-limit
Multicast Rate:      No-limit
Unknown Dest Mac Rate:      No-limit
QoS Profile:      QP3 configured by user
Ingress Rate Shaping :      Unsupported
Ingress IPTOS Examination:      Disabled
Ingress 802.1p Examination:      Enabled
Ingress 802.1p Inner Exam:      Disabled
Egress IPTOS Replacement:      Disabled
Egress 802.1p Replacement:      Disabled
NetLogin:      Enabled
NetLogin authentication mode:      MAC based
NetLogin port mode:      MAC based VLANs
Smart redundancy:      Enabled
Software redundant port:      Disabled
autopolarity:      Enabled
```

BlackDiamond 10808 and BlackDiamond 12800 series switches only. The following command displays more specific information for slot 1, port 1 on a BlackDiamond 10808 switch:

```
show ports 1:1 information detail
```

Following is sample output from this command:

```
Port: 1:1
  Virtual-router: VR-Default
  Type: UTP
  Random Early drop: Unsupported
  Admin state: Enabled with auto-speed sensing (100M Advertised), auto-
duplex (half-duplex Advertised)
  Link State: Active, 100Mbps, full-duplex
  Link Counter: Up 1 time(s)
  VLAN cfg:
    Name: peggy, Internal Tag = 4094, MAC-limit = No-limit

  STP cfg:

  Protocol:
    Name: peggy Protocol: ANY Match all protocols.
  Trunking: Load sharing is not enabled.
  EDP: Enabled
  ELSM: Disabled
  Learning: Enabled
  Unicast Flooding: Enabled
  Multicast Flooding: Enabled
  Broadcast Flooding: Enabled
  Jumbo: Disabled
  Link up/down SNMP trap filter setting: Enabled
  QoS Profile: None configured
  Aggregate Queue:
    QP0 MinBw = 0% PR = 100K Pri = 8
  Queue:
    QP1 MinBw = 0% MaxBw = 100% Pri = 1
    QP2 MinBw = 0% MaxBw = 100% Pri = 2
    QP3 MinBw = 0% MaxBw = 100% Pri = 3
    QP4 MinBw = 0% MaxBw = 100% Pri = 4
    QP5 MinBw = 0% MaxBw = 100% Pri = 5
    QP6 MinBw = 0% MaxBw = 100% Pri = 6
    QP7 MinBw = 0% MaxBw = 100% Pri = 7
    QP8 MinBw = 0% MaxBw = 100% Pri = 8
  Ingress Rate Shaping : support IQP1-2
    IQP1 MinBw = 0% MaxBw = 100% Pri = 1
    IQP2 MinBw = 0% MaxBw = 100% Pri = 2
  Ingress IPTOS Examination: Disabled
  Egress IPTOS Replacement: Disabled
  Egress 802.1p Replacement: Disabled
  NetLogIn: Disabled
  Smart redundancy: Enabled
  Software redundant port: Enabled
    Primary: 1:1
    Redundant: 1:2
    Redundant link configuration: Off
  autopolarity: Enabled
```

**NOTE**

On the BlackDiamond 10808 switch, the screen displays both ingress and egress QoS settings. The 10Gbps ports have 8 ingress queues, and the 1 Gbps ports have 2 ingress queues. (Refer to the ExtremeXOS Concepts Guide for more information on ingress queues, or bi-directional rate shaping.) The BlackDiamond 12800 series switch does not support ingress QoS settings.

BlackDiamond 12800 R-series switch only. The following command displays more specific information for slot 2, port 1 on a BlackDiamond 12800 R-series switch:

```
show ports 2:1 information detail
```

Following is sample output from this command:

```
Port: 2:1
Virtual-router: VR-Default
Type: NONE
Random Early drop: Unsupported
Admin state: Enabled with 10G full-duplex
Link State: Ready
Link Counter: Up 0 time(s)
VLAN cfg:
    Name: Default, Internal Tag = 1, MAC-limit = No-limit, Virtual ro
uter: VR-Default

STP cfg:
    s0(disable), Tag=(none), Mode=802.1D, State=FORWARDING

Protocol:
    Name: Default Protocol: ANY Match all protocols.
Trunking: Load sharing is not enabled.
EDP: Enabled
ELSM: Disabled
Learning: Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo: Disabled
High Priority COS: 7
Packet Length Adjustment: 0 bytes
QoS Profile: None configured
Aggregate Queue:
    QP0 MinBw = 0% MaxBw = 100% Pri = 8
Queue:
    QP1 MinBw = 0% MaxBw = 100% Pri = 1
    QP2 MinBw = 0% MaxBw = 100% Pri = 2
    QP3 MinBw = 0% MaxBw = 100% Pri = 3
    QP4 MinBw = 0% MaxBw = 100% Pri = 4
    QP5 MinBw = 0% MaxBw = 100% Pri = 5
    QP6 MinBw = 0% MaxBw = 100% Pri = 6
    QP7 MinBw = 0% MaxBw = 100% Pri = 7
    QP8 MinBw = 0% MaxBw = 100% Pri = 8
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
Egress IPTOS Replacement: Disabled
```

```

Egress 802.1p Replacement:    Disabled
NetLogin:                    Disabled
NetLogin port mode:          Port based VLANs
Smart redundancy:            Enabled
Software redundant port:     Disabled
Preferred medium:            Fiber

```

**NOTE**

On the GM-20XT and GM-20XTR modules for the BlackDiamond 12800 series switch and 12800 R-series switches, you can use each port as either copper or fiber; preferred medium displays which medium the specified port is using.

BlackDiamond 10808 switch and Summit X450a series switches (including SummitStack) with LW XENPAK ports only. The following command displays more specific information for the WAN PHY port on the LW XENPAK in port 1 of the Summit X450a -24t switch:

```

Port:    1
Framing: SONET
Clocking: Line
LoopBack: Off
Trace Section: alpha4
Trace Path:    delta6

```

History

This command was first available in ExtremeXOS 10.1.

Information on ingress rate shaping was added in ExtremeXOS 11.0.

NetLogIn, Smart Redundancy, and rate limiting were added in ExtremeXOS 11.1.

Information on unicast, multicast, and broadcast flooding; the `Port not present` parameter; and autopolarity status were added in ExtremeXOS 11.2.

The `netlogin` parameters were added in ExtremeXOS 11.3.

The BlackDiamond 12800 R-series specific parameters were added in ExtremeXOS 11.4.

Information on WAN PHY ports on the LW XENPAK modules on the BlackDiamond 10808 switch and the Summit X450a series switches was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms. (Information on WAN PHY ports on the LW XENPAK modules is available only on the BlackDiamond 10808 switch and the Summit X450a series switches, with or without SummitStack.)

show ports packet

```
show ports {mgmt | <port_list>} packet {no-refresh}
```

Description

Displays a snapshot or real-time histogram of packet statistics.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, the system displays information for all ports; if you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command. To clear the counters, use the `clear counters ports` command.

This status information may be useful for your technical support representative if you have a network problem.

The following packet statistics are displayed:

- Port number
- Link state
- Packet size

Example

The following command displays packet statistics for slot 1, port 1, slot 2, port 1, and slot 5, ports 1 through 8 on a modular switch:

```
show ports 1:1, 2:1, 5:1-5:8 packet
```

Following is sample output from this command:

```

Port      Link      Packet Sizes
      State  0-64    65-127  128-255  256-511  512-1023  1024-1518  Jumbo
=====
1:1  A      0        0        0        0        0        0        0
2:1  R      0        0        0        0        0        0        0
5:1  R      0        0        0        0        0        0        0
5:2  R      0        0        0        0        0        0        0
5:3  R      0        0        0        0        0        0        0
5:4  R      0        0        0        0        0        0        0
5:5  R      0        0        0        0        0        0        0
5:6  R      0        0        0        0        0        0        0
5:7  R      0        0        0        0        0        0        0
5:8  R      0        0        0        0        0        0        0
=====
Link Status: A-Active R-Ready, NP-Port not present, L-Loopback

```

History

This command was first available in ExtremeXOS 10.1.

The `Port Not Present` variable was added in ExtremeXOS 11.2.

The `no-refresh` variable was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show ports redundant

```
show ports redundant
```

Description

Displays detailed information about redundant ports.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information on software-controlled redundant ports on the switch:

```
show ports redundant
```

Following is sample output from this command:

```
Primary: *1:1          Redundant: 3:1, Link on/off option: OFF  
Flags: (*)Active, (!) Disabled, (g) Load Share Group
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show ports sharing

```
show ports sharing
```

Description

Displays port load-sharing groups, or link aggregation groups (LAGs).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Output from this command displays the following information:

- Config Master—The port that is configured as the master logical port of the link aggregation group (LAG). This number is also the LAG group ID.
- Current Master—In LACP, this is the port that is currently the LAG group ID, or master logical port for the LAG.
- Agg Control—This is the aggregation control for the specified LAG; it can be either static or LACP. In LACP, it is the aggregation control for the specified LAG.
- Ld Share Algorithm—The algorithm used for the link aggregation. The available link aggregation algorithms vary among platforms; see the *ExtremeXOS Concepts Guide* for more information.
- Ld Share Group—The specific ports that belong to each LAG, or the port numbers in the trunk. A port can belong to only one LAG, either static or dynamic.
- Agg Mbr—In LACP, this shows whether the port has been added to the aggregator or not; it will be either Y for yes or - for no.
- Link Status—This is the current status of the link
- Link Up transitions—Number of times the link has cycled through being up, then down, then up.

Example

The following command displays link aggregation on the switch:

```
show ports sharing
```

The following is sample output from this command on the Summit X450 series switch:

```
Load Sharing Monitor
Config      Current      Agg      Ld Share      Ld Share      Agg      Link      Link Up
Master      Master      Control  Algorithm      Group      Mbr      Status    transitions
=====
      5        5        LACP      L2             5          Y         A         3
              L2             6          Y         A         3
              L2             7          Y         A         3
              L2             8          Y         A         3
              L2             9          -         A         3
              L2            10         -         A         3
      12              Static    L2            12         -         R         0
              L2            13         -         R         0
=====
Link Status: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
Default algorithm: L2
Number of load sharing trunks: 2
```

History

This command was first available in ExtremeXOS 10.1.

The LACP (dynamic link aggregation) feature was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show ports utilization

```
show ports {mgmt | <port_list> | stack-ports <stacking-port-list>}
utilization {bandwidth | bytes | packets}
```

Description

Displays real-time port utilization information. The total utilization displays as real-time information, constantly refreshing, and the parameter displays show a snapshot of the activity on the port when you issue the command.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
stacking-port-list	Specifies one or more stacking slots and ports.
bandwidth	Specifies port utilization as percentage of bandwidth.
bytes	Specifies port utilization in bytes per second.
packets	Specifies port utilization in packets per second.

Default

N/A.

Usage Guidelines

The software continuously monitors port utilization and calculates bandwidth as a function of each port's maximum link capacity.

The total utilization display presents real-time statistics. Use the <spacebar> to toggle the real-time displayed information for packets, bytes, and bandwidth in that order. When you use a parameter (packets, byte, or bandwidth) with the command, the display for the specified type shows a snapshot per port when you issued the command.

If you do not specify a port number or range of ports, port utilization information is displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays utilization statistics for port 1 on a stand-alone switch:

```
show ports 1 utilization
```

The following command displays utilization statistics for slot 3, port 1 on a modular switch:

```
show ports 3:1 utilization
```

The following example shows sample output from the `show ports utilization` command for all ports on a Summit X450 series switch that is not part of a SummitStack (real-time display that constantly refreshes):

```

Link Utilization Averages
Port      Link      Receive      Peak Rx      Transmit      Peak Transmit
          Status    packets/sec  packets/sec  packets/sec   packets/sec
=====
1          A          0            0            0            0
2          R          0            0            0            0
3          R          0            0            0            0
4          R          0            0            0            0
5          R          0            0            0            0
6          R          0            0            0            0
7          R          0            0            0            0
8          R          0            0            0            0
9          R          0            0            0            0
10         R          0            0            0            0
11         R          0            0            0            0
12         R          0            0            0            0
13         R          0            0            0            0
14         R          0            0            0            0
15         R          0            0            0            0
16         R          0            0            0            0
17         R          0            0            0            0
=====
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up  D->page down ESC->exit

```



NOTE

Use the `<spacebar>` to toggle this real-time display for all ports from packets to bytes and bandwidth, in that order.

The following example shows sample output from the `show ports utilization bandwidth` command:

```

Port      Link      Link      Receive      Peak Rx      Transmit      Peak Transmit
          Status    Speed    % bandwidth  % bandwidth  % bandwidth  % bandwidth
=====
1          A      1000      0.00          0.00          0.00          0.00
2          R          0        0.00          0.00          0.00          0.00
3          R          0        0.00          0.00          0.00          0.00
4          R          0        0.00          0.00          0.00          0.00
5          R          0        0.00          0.00          0.00          0.00
6          R          0        0.00          0.00          0.00          0.00
7          R          0        0.00          0.00          0.00          0.00
8          R          0        0.00          0.00          0.00          0.00
9          R          0        0.00          0.00          0.00          0.00
10         R          0        0.00          0.00          0.00          0.00
11         R          0        0.00          0.00          0.00          0.00
12         R          0        0.00          0.00          0.00          0.00
13         R          0        0.00          0.00          0.00          0.00
14         R          0        0.00          0.00          0.00          0.00
15         R          0        0.00          0.00          0.00          0.00
16         R          0        0.00          0.00          0.00          0.00
17         R          0        0.00          0.00          0.00          0.00

```

```

18      R      0      0.00      0.00      0.00      0.00
19      R      0      0.00      0.00      0.00      0.00
20      R      0      0.00      0.00      0.00      0.00
21      R      0      0.00      0.00      0.00      0.00
22      R      0      0.00      0.00      0.00      0.00
23      R      0      0.00      0.00      0.00      0.00
24      R      0      0.00      0.00      0.00      0.00
25      R      0      0.00      0.00      0.00      0.00
26      R      0      0.00      0.00      0.00      0.00
=====

```

Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback

The following example shows sample output of the show ports utilization bytes command:

```

Port      Link      Receive      Peak Rx      Transmit      Peak Transmit
      Status bytes/sec bytes/sec bytes/sec bytes/sec
=====
1          A          0          0          0          0
2          R          0          0          0          0
3          R          0          0          0          0
4          R          0          0          0          0
5          R          0          0          0          0
6          R          0          0          0          0
7          R          0          0          0          0
8          R          0          0          0          0
9          R          0          0          0          0
10         R          0          0          0          0
11         R          0          0          0          0
12         R          0          0          0          0
13         R          0          0          0          0
14         R          0          0          0          0
15         R          0          0          0          0
16         R          0          0          0          0
17         R          0          0          0          0
18         R          0          0          0          0
19         R          0          0          0          0
20         R          0          0          0          0
21         R          0          0          0          0
22         R          0          0          0          0
23         R          0          0          0          0
24         R          0          0          0          0
25         R          0          0          0          0
26         R          0          0          0          0
=====

```

Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback

The following example shows sample output of the `show ports utilization packets` command:

Port	Link Status	Receive packets/sec	Peak Rx packets/sec	Transmit packets/sec	Peak Transmit packets/sec
=====					
1	A	0	0	0	0
2	R	0	0	0	0
3	R	0	0	0	0
4	R	0	0	0	0
5	R	0	0	0	0
6	R	0	0	0	0
7	R	0	0	0	0
8	R	0	0	0	0
9	R	0	0	0	0
10	R	0	0	0	0
11	R	0	0	0	0
12	R	0	0	0	0
13	R	0	0	0	0
14	R	0	0	0	0
15	R	0	0	0	0
16	R	0	0	0	0
17	R	0	0	0	0
18	R	0	0	0	0
19	R	0	0	0	0
20	R	0	0	0	0
21	R	0	0	0	0
22	R	0	0	0	0
23	R	0	0	0	0
24	R	0	0	0	0
25	R	0	0	0	0
26	R	0	0	0	0
=====					
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback					

The following example shows sample output of the `show ports stack-ports utilization` command:

Port	Link	Receive	Peak Rx	Transmit	Peak Transmit
Status		packets/sec	packets/sec	packets/sec	packets/sec
=====					
1:1	A	267	267	367	367
1:2	A	501	501	897	897
2:1	A	249	249	574	574
2:2	A	897	897	501	501
3:1	A	8	8	240	240
3:2	A	574	574	249	249
4:1	A	369	369	267	267
4:2	A	30	30	29	29
5:1	A	30	30	31	31
5:2	A	25	25	25	25
6:1	A	25	25	25	25
6:2	A	11	11	10	10
7:1	A	119	119	3	3
7:2	A	10	10	11	11
8:1	A	3	3	119	119

```
8:2 A 240 240 8 8
```

```
=====
```

```
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms. The stack-ports option is available only on SummitStack and the Summit Family of switches.

show ports wan-phy configuration

```
show ports <port_list> wan-phy configuration
```

Description

Displays the configuration of the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.

Default

Real-time information.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

This command displays the configuration for the specified WAN PHY ports. This command displays:

- Port number
- Link
- Framing
- Clocking
- Loopback
- Section trace
- Path trace

If you do not specify a port number or range of ports, the configuration is displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Example

The following command displays the configuration for all LW XENPAK WAN PHY ports:

```
show ports wan-phy configuration
```

Following is sample output from this command:

```

WAN PHY Port Configuration
Port Link Framing Clocking LoopBack  SectionTrace  PathTrace
=====
1   A   SONET  Internal  off      abcdefghijklmn  01234567890
2   R   SONET  Internal  Line    abcdefghijklmn  01234567890
3   R   SONET  Line      off      abcdefghijklmn  01234567890
4   A   SONET  Internal  Line    abcdefghijklmn  01234567890
5   R   SONET  Line      off      abcdefghijklmn  01234567890
6   A   SONET  Internal  Line    abcdefghijklmn  01234567890
7   R   SDH    Line      Line    abcdefghijklmn  01234567890
=====

```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switch (whether or not included in SummitStack), and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

show ports wan-phy errors

```
show ports <port_list> wan-phy errors {no-refresh}
```

Description

Displays the error information of the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

This command displays the counters associated with parity errors for the specified WAN PHY ports. The display shows the following:

- Port number
- B1, B2, and B3
- Far end path block error count
- Far end line BIP error count

If you do not specify a port number or range of ports, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Example

The following command displays the errors for the LW XENPAK WAN PHY ports 1 to 4:

```
show ports 1-4 wan-phy errors
```

Following is sample output from this command:

```

WAN PHY Port Errors
Port          B1          B2          B3          FELBE          FEPL
=====
1            0123456789 01234567890123456789 0123456789 0123456789
2            0123456789 01234567890123456789 0123456789 0123456789
3            0123456789 01234567890123456789 0123456789 0123456789
4            0123456789 01234567890123456789 0123456789 0123456789

=====
0->Clear Counters  U->page up  D->page down ESC->exit
FELBE = Far End Path Block Error Count
FEPBE = Far End Line BIP Error Count

```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switch (whether or not included in a SummitStack), and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

show ports wan-phy events

```
show ports <port_list> wan-phy events {no-refresh}
```

Description

Displays the events information of the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

This command displays the events for the specified WAN PHY ports. The display shows whether each parameter is in the alarm or normal state; the alarm state includes alarm, warning, and error states.

This command displays notices for the following events:

- LOS: loss of signal
- SEF: Severely Error Frame
- LOF: loss of frame
- LOP: loss of pointer
- AIS-L: alarm indication signal, line
- AIS-P: alarm indication signal, path
- RDI-L: remote defection indication, line
- FarEnd PLM-P/LCD-P: Far end path label mismatch / loss of code-group delineation
- FarEnd AIS-P/LCD-P: Far end path alarm indication signal / path loss of pointer
- PLM-P: Path label mismatch
- LCD-P: path loss of code-group delineation

If you do not specify a port number or range of ports, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Example

The following command displays the events for the LW XENPAK WAN PHY ports 1 to 4:

```
show ports 1-4 wan-phy events
```

Following is sample output from this command:

```
WAN PHY Port Events:
Port LOS SEF LOF LOP AIS-L AIS-P RDI-L FarEnd FarEnd PLM-P LCD-P
      PLM-P/LCD-P AIS-P/LOP-P
=====
25   A   N   A   N   N   N   N   N           N       N       A
26   N   N   A   N   N   N   N   A           N       A       N
=====
Event Status: A-Alarm N-Normal
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switch (whether or not included in a SummitStack), and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

show ports wan-phy overhead

```
show ports <port_list> wan-phy overhead {no-refresh}
```

Description

Displays selected WAN PHY OAM overhead information for the specified WAN PHY port.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

This command displays selected WAN PHY OAM overhead for the specified WAN PHY ports. The overhead number comes directly from the WAN PHY framing; for each overhead, both transmit and receive numbers are displayed.



NOTE

Overhead values display in hexadecimal format.

If you do not specify a port number or range of ports, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Example

The following command displays the events for the LW XENPAK WAN PHY ports 1 to 3:

```
show ports 1-3 wan-phy overhead
```

Following is sample output from this command:

```

WAN PHY Port Overhead Info
Port      C2      G1      K1      K2      M1      J0(0-15)      J1(0-15)
          Rx      Rx      Rx      Rx      Rx      Rx              Rx
=====
25         0       0       0       0       F8  abcdefghijklmn  ExtremeNetworks
26        1A       0       0       0       F8  ExtremeNetworks  abcdefghijklmn
=====
Overhead values for C2, G1, K1, K2, and M1 are displayed in hexadecimal
representation.
Overhead values for J0 and J1 are displayed in ascii alphanumeric
representation with a dot representing a non-alphanumeric character.

```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switch (whether or not included in a SummitStack), and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

show slot

```
show slot {<slot> {detail} | detail }
```

Description

Displays the slot-specific information.

Syntax Description

slot	Specifies a slot on a modular switch or SummitStack.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines



NOTE

You may see slightly different information displayed depending on the platform and configuration you are using.

The `show slot` command displays the following information:

- The slot number
- The type of module installed in the slot
- The type of module configured for the slot
- The state of the module, whether the power is down, if the module is operational, if a diagnostic being run, if there is a mismatch between the slot configuration and the module in the slot
- The number of ports on the module

If you do not specify a slot number, information for all slots is displayed.

The display also includes a notice of insufficient power, should that arise.

The `show slot` command displays the following states, among others:

- Empty (BD 8800 series switch only—This also displays if you have a module in the chassis that is unsupported by the current software you are running.)
- Down
- Power ON
- Powered OFF
- Booting
- Initializing
- VLAN sync

- FDB sync
- ACL sync
- RT sync
- Operational

Show slot command on a SummitStack. The output of the show slot command is different for a SummitStack. The output of this command displays eight rows. If a node in the Active Topology is not assigned a slot number, the state of the slot is shown as Empty. A node that shows as a slot has successfully joined the active topology. This means, EXOS can communicate with this node and does not mean that the node has been successfully brought up. The card state in the display indicates whether the slot was successfully started. If the card state is Operational, then the node is being used in the stack to carry user data as configured.

The output of the command displays the slot number, type of the Summit in that slot, state of the slot, and the number of ports.

The number of ports does not include the stacking links. It includes the option card ports regardless of whether the option card is installed.

This command is not available on Summit family switches operating in non stacking mode.

Example

You see slightly different displays, depending on the platform.

Switches on a SummitStack. The following example displays module information for all slots in a stack:

```
* Slot-7 Stack.1 # show slot
```

Slots	Type	Configured	State	Ports
Slot-1	SummitX450-24x	SummitX450-24x	Operational	26
Slot-2	X450a-48t	X450a-48t	Operational	50
Slot-3	X450a-24t	X450a-24t	Operational	26
Slot-4	SummitX450-24x	SummitX450-24x	Operational	26
Slot-5	X450e-24p	X450e-24p	Operational	26
Slot-6	X450e-24p	X450e-24p	Operational	26
Slot-7	X450e-24p	X450e-24p	Operational	26
Slot-8			Empty	0

The following example displays module information for a specific slot on the stack:

```
* Slot-7 Stack.91 # show slot 1 detail
```

Slot-1 information:

State:	Operational
Download %:	100
Restart count:	4
Serial number:	800187-00-02 0635G-00074
Hw Module Type:	SummitX450-24x
SW Version:	12.0.0.17
SW Build:	v1170b17
Configured Type:	SummitX450-24x
Ports available:	26
Recovery Mode:	Reset
Node MAC:	02:04:96:27:87:17
Current State:	STANDBY

```

Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg

```

The following example displays detailed module information for all slots on a stack:

```

* Slot-7 Stack.90 # show slot detail
Slot-1 information:
  State:                Operational
  Download %:           100
  Restart count:        4
  Serial number:        800187-00-02 0635G-00074
  Hw Module Type:       SummitX450-24x
  SW Version:           12.0.0.17
  SW Build:             v1170b17
  Configured Type:      SummitX450-24x
  Ports available:      26
  Recovery Mode:        Reset
  Node MAC:             02:04:96:27:87:17
  Current State:        STANDBY
  Image Selected:       secondary
  Image Booted:         secondary
  Primary ver:          12.0.0.16
  Secondary ver:        12.0.0.17
  Config Selected:      primary.cfg

Slot-2 information:
  State:                Operational
  Download %:           100
  Restart count:        4
  Serial number:        800163-00-04 0635G-01187
  Hw Module Type:       X450a-48t
  SW Version:           12.0.0.17
  SW Build:             v1170b17
  Configured Type:      X450a-48t
  Ports available:      50
  Recovery Mode:        Reset
  Node MAC:             02:04:96:27:87:17
  Current State:        STANDBY
  Image Selected:       secondary
  Image Booted:         secondary
  Primary ver:          12.0.0.16
  Secondary ver:        12.0.0.17
  Config Selected:      primary.cfg

Slot-3 information:
  State:                Operational
  Download %:           100
  Restart count:        4
  Serial number:        800152-00-04 0630G-00736
  Hw Module Type:       X450a-24t
  SW Version:           12.0.0.17
  SW Build:             v1170b17

```

```

Configured Type:      X450a-24t
Ports available:      26
Recovery Mode:        Reset
Node MAC:              02:04:96:27:87:17
Current State:         STANDBY
Image Selected:        secondary
Image Booted:          secondary
Primary ver:           12.0.0.16
Secondary ver:          12.0.0.17
Config Selected:       primary.cfg

```

Slot-4 information:

```

State:                 Operational
Download %:            100
Restart count:         4
Serial number:          0635G-00073 S450-24X
Hw Module Type:        SummitX450-24x
SW Version:             12.0.0.17
SW Build:               v1170b17
Configured Type:        SummitX450-24x
Ports available:        26
Recovery Mode:          Reset
Node MAC:               02:04:96:27:87:17
Current State:          STANDBY
Image Selected:         secondary
Image Booted:           secondary
Primary ver:            12.0.0.16
Secondary ver:           12.0.0.17
Config Selected:        primary.cfg

```

Slot-5 information:

```

State:                 Operational
Download %:            100
Restart count:         4
Serial number:          800153-00-04 0646G-00683
Hw Module Type:        X450e-24p
SW Version:             12.0.0.17
SW Build:               v1170b17
Configured Type:        X450e-24p
Ports available:        26
Recovery Mode:          Reset
Node MAC:               02:04:96:27:87:17
Current State:          STANDBY
Image Selected:         secondary
Image Booted:           secondary
Primary ver:            12.0.0.16
Secondary ver:           12.0.0.17
Config Selected:        primary.cfg

```

Slot-6 information:

```

State:                 Operational
Download %:            100
Restart count:         4
Serial number:          800153-00-04 0646G-00691
Hw Module Type:        X450e-24p
SW Version:             12.0.0.17

```

```

SW Build:          v1170b17
Configured Type:   X450e-24p
Ports available:   26
Recovery Mode:     Reset
Node MAC:          02:04:96:27:87:17
Current State:     BACKUP
Image Selected:    secondary
Image Booted:      secondary
Primary ver:       12.0.0.16
Secondary ver:     12.0.0.17
Config Selected:   primary.cfg

```

Slot-7 information:

```

State:             Operational
Download %:        100
Restart count:     4
Serial number:     800153-00-01 0603G-00741
Hw Module Type:    X450e-24p
SW Version:        12.0.0.17
SW Build:          v1170b17
Configured Type:   X450e-24p
Ports available:   26
Recovery Mode:     Reset
Node MAC:          02:04:96:27:87:17
Current State:     MASTER
Image Selected:    secondary
Image Booted:      secondary
Primary ver:       12.0.0.16
Secondary ver:     12.0.0.17
Config Selected:   primary.cfg

```

Slot-8 information:

```

State:             Empty
Restart count:     4
Serial number:
Hw Module Type:
Configured Type:
Ports available:   0
Recovery Mode:     Reset
Node MAC:          00:00:00:00:00:00
Current State:
Image Selected:
Image Booted:
Primary ver:
Secondary ver:
Config Selected:

```

BlackDiamond 8800 series switch only. The following example displays module information for all slots:

Slots	Type	Configured	State	Ports	Flags
Slot-1			Empty	0	
Slot-2	G24X	G24X	Operational	24	M S
Slot-3			Empty	0	
Slot-4			Empty	0	
Slot-5	G8X	G8X	Operational	8	M S

Slot-6		Empty	0		
Slot-7	G48P	Empty	48		
Slot-8	G48P	Operational	48	M	S
Slot-9	10G4X	Powered OFF	4		SI
Slot-10		Empty	0		
MSM-A	MSM-G8X	Operational	0		S
MSM-B		Empty	0		

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")

The following example displays module information for a specified slot on the BlackDiamond 8810 switch:

```
BD-8810.3 # show slot 2
```

```
Slot-2 information:
  State:                Operational
  Download %:           100
  Flags:                M S
  Serial number:        800115-00-02 04314-00006
  Hw Module Type:       G24X
  SW Version:           11.1.0.0
  SW Build:             v11100 by Wed Oct 20 16:01:52 PDT 2004
  Configured Type:      G24X
  Ports available:      24
```

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")

BlackDiamond 10808 and BlackDiamond 12800 series switches only. The following example displays module information for all slots:

```
BD-10808.2 # show slot
```

Slots	Type	Configured	State	Ports	Flags
Slot-1			Empty	0	
Slot-2			Empty	0	
Slot-3			Empty	0	
Slot-4			Empty	0	
Slot-5			Empty	0	
Slot-6			Empty	0	
Slot-7	G60T		Operational	60	M
Slot-8			Empty	0	
MSM-A	MSM-1XL		Operational	0	
MSM-B	MSM-1XL		Present	0	

```
Flags : M - Backplane link to Master MSM is Active
        B - Backplane link to Backup MSM is also Active
        D - Slot Disabled, S - Slot Secured
        I - Insufficient Power (refer to "show power budget")
        E - Errors on slot will be ignored (no corrective action initiated)
```

The following example displays module information for a specified slot on the BlackDiamond 10808 switch:

```
BD-10808.3 # show slot 7
```

```
Slot-7 information:
```

```
State:           Operationa
Download %:      100
Flags:          M
Serial number:   804403-00-02 0330Z-00086
Hw Module Type:  G60T
SW Version:      11.1.1.1
SW Build:        v1111b1 by release-manager on Fri Nov 5 06:42:22 PST 2
```

```
004
```

```
Configured Type: G60T
Ports available:  60
```

```
Flags : M - Backplane link to Master MSM is Active
        B - Backplane link to Backup MSM is also Active
        D - Slot Disabled, S - Slot Secured
        I - Insufficient Power (refer to "show power budget")
        E - Errors on slot will be ignored (no corrective action initiated)
```

You see the same output if you add the optional parameter `detail` to the command for displaying a specified slot number.

The following example displays detailed module information for all slots on a BlackDiamond 10808 switch:

```
BD-10808.4 # show slot detail
```

```
Slot-1 information:
```

```
State:           Operational
Download %:      100
Flags:          MB S
```

```

Serial number:      804403-00-08 0405F-00063
Hw Module Type:    G60T
SW Version:        11.4.0.23
SW Build:          v1140b23 by release-manager on Thu Feb 16 11:43:35 PST 2
006
Configured Type:    G60T
Ports available:    60

Slot-2 information:
State:              Powered OFF
Download %:         0
Flags:              I
Last Error:         Insufficient power
Serial number:      804406-00-02 0511F-00605
Hw Module Type:    10G2H
Configured Type:
Ports available:    42

Slot-3 information:
State:              Powered OFF
Download %:         0
Flags:              I
Last Error:         Insufficient power
Serial number:      804406-00-02 0511F-00602
Hw Module Type:    10G2H
Configured Type:
Ports available:    42

Slot-4 information:
State:              Empty
Flags:
Serial number:
Hw Module Type:
Configured Type:
Ports available:    0

Slot-5 information:
State:              Empty
Flags:
Serial number:
Hw Module Type:
Configured Type:
Ports available:    0

Slot-6 information:
State:              Empty
Flags:
Serial number:
Hw Module Type:
Configured Type:
Ports available:    0

Slot-7 information:
State:              Empty
Flags:
Serial number:

```



```

    Hw Module Type:
    Configured Type:
    Ports available:      0

Slot-8 information:
    State:                Empty
    Flags:
    Serial number:
    Hw Module Type:
    Configured Type:
    Ports available:      0

MSM-A information:
    State:                Operational
    Download %:           100
    Flags:                S
    Serial number:        804301-00-09 0434F-00125
    Hw Module Type:       MSM-1XL
    SW Version:           11.4.0.26
    Configured Type:
    Ports available:      0

MSM-B information:
    State:                Operational
    Download %:           100
    Flags:                S
    Last Error:           Lost Connectivity to MSM
    Serial number:        804301-00-09 0437F-00026
    Hw Module Type:       MSM-1XL
    Configured Type:
    Ports available:      0

Flags : M - Backplane link to Master MSM is Active
        B - Backplane link to Backup MSM is also Active
        D - Slot Disabled, S - Slot Secured
        I - Insufficient Power (refer to "show power budget")

```

History

This command was first available in ExtremeXOS 10.1.

This command was first available on SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

unconfigure ports display string

```
unconfigure ports <port_list> display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

This command removes the display string that you configured using the `configure ports display-string` command.

Example

The following command clears the user-defined display string from slot 2, port 4 on a modular switch:

```
unconfigure ports 2:4 display-string
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure ports redundant

```
unconfigure ports <port_list> redundant
```

Description

Clears a previously configured software-controlled redundant port.

Syntax Description

port_list	This refers to the primary port of the redundant pair and specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The list of port numbers or the port display string specifies the primary port(s).

Example

The following command unconfigures a software-controlled redundant port on a modular switch:

```
unconfigure ports 2:3 redundant
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure ports wan-phy

```
unconfigure ports [<port_list> | all] wan-phy
```

Description

Resets the configuration parameters of the specified WAN PHY port to default values.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
all	Specified all WAN PHY ports.

Default

N/A.

Usage Guidelines

This command is visible and available only if you have an LW XENPAK in your system.

The default configurable WAN PHY OAM parameters are as follows:

- Framing—SONET
- Clock source—line
- Section trace—15-character NULL string
- Path trace—15-character NULL string
- Loopback—off

Example

The following command unconfigures the configurable WAN PHY OAM parameters on all LW XENPAK ports:

```
unconfigure ports all wan-phy
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on LW XENPAK ports on the BlackDiamond 10808 switch, the Summit X450a series switch (whether or not included in a SummitStack), and beginning in ExtremeXOS 12.0 on the BlackDiamond 12800 series switch.

unconfigure slot

```
unconfigure slot <slot>
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

slot	Specifies a slot on a modular switch.

Default

N/A.

Usage Guidelines

If you issue the `unconfigure ports wan-phy` command on a slot containing a module with any ports configured for software-controlled redundancy, this command wipes away all software-controlled redundancy on both ports; both ports return to normal. Refer to *ExtremeXOS Concepts Guide* for more information on software-controlled redundant ports.

Example

The following command clears slot 4 of a previously assigned module type:

```
unconfigure slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

6

Universal Port Management Commands

This chapter describes commands for:

- Configuring a Universal Port via CLI

The Universal Port is a flexible framework that allows the switch to take direct action based on events.

The CLI-based scripting capability in ExtremeXOS allows users to create dynamic profiles that are activated by a trigger event, and cause dynamic changes to the configuration of the switch to enforce a policy.

**NOTE**

Beginning in ExtremeXOS Release 12.0, the Universal Port feature is available with an Edge license. In earlier releases, an Advanced Edge license is required.

configure upm event

```
configure upm event <upm-event> profile <profile-name> ports <port_list>]
```

Description

Configures a pre-defined event that triggers the named profile.

Syntax Description

<upm-event>	Specifies a pre-defined event type: device-detect, device-undetected, user-authenticate, user-unauthenticated
profile <profile-name>	Specifies the profile to be configured.
ports <port-list>	Attaches the UPM profile to the specified port(s).

Default

N/A.

Usage Guidelines

This command configures a profile to be executed when the specified event occurs on the specified port(s).

You can configure multiple user profiles on the same port(s).

Example

The following example shows how to configure a profile on port 1:1, called “profile 1” that is triggered by the event “device-detect”:

```
* BD-10808.35 # configure upm event device-detect profile "p1" ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

configure upm timer after

```
configure upm timer <timer-name> after <time-in-secs> {every <seconds>}
```

Description

Creates and names a UPM timer that will be activated after the specified time in seconds.

Syntax Description

<timer-name>	Specifies the name of the UPM timer to be created.
<time-in-secs>	Configures the interval after which the UPM timer will be activated.
every <seconds>	Configures the UPM timer to be activated after every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to configure a timer that activates after the specified time. This is useful for deployment in CLI scripts, because you do not know what the current time will be when the script executes.

When a switch configuration is saved or restored, the UPM timers will be activated only at the predetermined timings that were originally configured with the start time.

The periodic timer configured with the `every` keyword and the one-time timer configured with only the `after` keyword have a maximum range of one year in seconds (31,622,400 seconds).

Example

The following example configures the UPM timer "A" to be activated every 10 seconds, after an interval of 20 seconds:

```
* BD-10808.24 # configure upm timer "timerA" after 20 every 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure upm timer at

```
configure upm timer <timer-name> at <month> <day> <year> <hour> <min>
<secs> {every <seconds>}
```

Description

Use this command to configure the time setting on a UPM timer.

Syntax Description

<timer-name>	Specifies the name of the UPM timer to be created.
<month>	Configures the month when the UPM timer will be activated.
<day>	Configures the day when the UPM timer will be activated.
<year>	Configures the year when the UPM timer will be activated.
<hour>	Configures the hour when the UPM timer will be activated.
<min>	Configures the minute when the UPM timer will be activated.
<secs>	Configures the second when the UPM timer will be activated.
every <seconds>	Configures the UPM timer to be activated at every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to when you know the exact time you want an event to execute. If you use this command without the `every` keyword, the timer will be activated once at the specified time. The `every` keyword configures a periodic timer that will be activated at every instance of the time specified in seconds.

When a switch configuration is saved or restored, the UPM timers will be activated only at the predetermined timings that were originally configured with the start time.

Example

The following example shows how to configure a timer, "T1" that will be activated every 10 seconds beginning at 1400 hours on October 16, 2006:

```
* BD-10808.33 # configure upm timer "t1" at 10 16 2006 14 00 00 every 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure upm timer profile

```
configure upm timer <timerName> profile <profileName>
```

Description

Associates a profile with a UPM timer.

Syntax Description

<timer-name>	Specifies the name of the UPM timer to be associated with the named profile.
<profileName>	Specifies the name of the profile to be associated with the UPM timer.

Default

N/A.

Usage Guidelines

Each timer can be attached to only one profile. Once a timer is configured to a profile, it must be unconfigured from that profile before it can be configured to a different profile.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

create upm profile

```
create upm profile <profile-name>
```

Description

Creates a new profile of a specified type.

Syntax Description

<profile-name>	Specifies the UPM profile to be created.

Default

N/A.

Usage Guidelines

Use this command to create a profile and name it. The maximum profile size is 5000 characters.

There is a limited capability to edit the profile with this command. If you enter a period (.) as the first and the only character on a line, you will terminate the editing of the file. Use the command: `edit upm profile <profile-name>` for block mode capability.

Example

The following example shows how to create a profile named "P2":

```
*BD-10808.22 # create upm profile p2
enab por 2:
dis por 3:1
.
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

create upm timer

```
create upm timer <timer-name>
```

Description

Creates and names a UPM timer.

Syntax Description

<timer-name>	Specifies the name of the UPM timer to be created.

Default

N/A.

Usage Guidelines

You can create UPM timers with a name. A profile can be associated with eight timers, but a timer can be bound to only one profile at any point in time. You can create a maximum of 32 timers. A name space for the timers is available to help when you are typing the commands.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete upm profile

```
delete upm profile <profile-name>
```

Description

Deletes the specified profile.

Syntax Description

<profile-name>	Specifies the UPM profile to be deleted.

Default

N/A.

Example

The following command deletes a UPM profile called *sample_1*:

```
delete upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete upm timer

```
delete upm timer <timer-name>
```

Description

Deletes the specified UPM timer.

Syntax Description

<timer-name>	Specifies the name of the UPM timer to be deleted.

Default

N/A.

Usage Guidelines

You can delete a UPM timer by specifying its name.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable upm profile

```
disable upm profile <profile-name>
```

Description

Disables the use of CLI scripting on the switch.

Syntax Description

<profile-name>	Specifies the UPM profile to be disabled.

Default

A UPM profile is enabled by default.

Example

The following command disables a UPM profile called *sample_1* on the switch:

```
disable upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

edit upm profile

```
edit upm profile <profile-name>
```

Description

Allows you to edit the specified profile.

Syntax Description

<profile-name>	Specifies the UPM profile to be edited.

Default

N/A.

Usage Guidelines

Use the command to have VI-like editor features for editing the profile. Changes appear when you close the file for editing, not when you save it.

History

This command was first available in ExtremeXOS 11.6.

enable upm profile

```
enable upm profile <profile-name>
```

Description

Enables the use of CLI scripting on the switch.

Syntax Description

<profile-name>	Specifies the UPM profile to be enabled.

Default

A UPM profile is enabled by default.

Example

The following command enables a UPM profile called *example* on the switch:

```
enable upm profile example
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show upm events

```
show upm events <event-type>
```

Description

Displays UPM events of the specified type.

Syntax Description

<event-type>	Displays events of the specified type for all profiles. Valid values for event-type are:
	• device-detect • device-remove • user-authenticate • user-unauthenticated

Default

N/A.

Usage Guidelines

Use this command to display the following types of events:

- device-detect
- device-remove
- user-authenticate
- user-unauthenticated

Example

The following command displays device-detect events:

```
show upm event device-detect
```

The output of the command is similar to the following:

```
-----
UPM Profile           PortList
-----
profile1              3
-----
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show upm history

```
show upm history {profile <profile-name> | exec-id <number> | event <upm-
event> | status [pass | fail] | timer <timer-name> | detail}
```

Description

Displays (in a tabular column) the history of a UPM profile executed on the switch.

Syntax Description

<profile-name>	Displays the execution history of the specified profile.
exec-id <number>	Displays the contents of a specific profile execution.
event <upm-event>	Displays UPM events that have executed for the specified profile.
status	Displays the execution status.
pass	Indicates the execution was successful.
fail	Indicates the execution failed.
timer <timer-name>	Filters the display by timers.
detail	Displays the complete history of the profile.

Default

N/A.

Usage Guidelines

This is useful for trouble shooting and testing

Example

The output of the command is similar to the following:

```
* BD-10808.41 # show upm history exec-id 51
UPM Profile: p1
Event: UPM-Timer-t1 , Time run: 2006-10-16 14:02:10
Execution Identifier: 51 Execution Status: Pass
Execution Information:
1 # enable cli scripting
2 # configure cli mode non-persistent
3 # set var EVENT.NAME TIMER-AT
4 # set var EVENT.TIMER_TYPE PERIODIC
5 # set var EVENT.TIME 1161007330
6 # set var EVENT.TIMER_DELTA_SECS 0
7 # set var EVENT.PROFILE p1
8 # ena por 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show upm profile

```
show upm profile <name>
```

Description

Displays a list of the UPM profiles on the system and some of their configuration information, or the contents of a specified profile.

Syntax Description

<name>	Displays the contents of the specified profile.

Default

N/A.

Usage Guidelines

To see a list of all UPM profiles on a switch, use the command without the `name` option. The resulting display shows the names of the profiles on the system and their status, active or disabled.

Use the `name` option to see the contents of a specific profile.

Example

The output of the command is similar to the following:

```
* BD-10808.36 # show upm profile
=====
UPM Profile      Events              Ports              Flags
=====
p1               UPM Timer(t1)
p1               device detect      1:1                e
p2
=====
Number of UPM Profiles: 2
Flags: d - disabled, e - enabled

* BD-10808.37 # show upm profile "p1"
Created at : 2010-04-11 04:07:41
Last edited at : 2010-04-11 04:07:41

*****Profile Contents Begin*****
ena por 1:1

*****Profile Contents Ends*****

Profile State: Enabled
```



```

Events and ports configured on the profile:
=====
Event                               Port list
=====
device detect                       :    1:1
device undetect                     :
User Authenticated                   :
User Unauthenticated                :
=====
* BD-10808.38 #

* BD-10808.37 # show upm profile "p1"
Created at : 2010-04-11 04:07:41
Last edited at : 2010-04-11 04:07:41
*****Profile Contents Begin*****
ena por 1:1
*****Profile Contents Ends*****
Profile State: Enabled
Events and ports configured on the profile:
=====
Event Port list
=====
device detect : 1:1
device undetect :
User Authenticated :
User Unauthenticated :
=====

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show upm timers

```
show upm timers
```

Description

Displays a list of the UPM timers on the system and some of their configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to see a list of all UPM timers on a switch. The resulting display shows timer configurations, the associated profile, and flags to indicate the timer status. Flags are defined below:

- a - active
- d - disabled
- p - periodic

Example

This command displays UPM timer configuration:

```
show upm timers
```

The output of this command is similar to the following:

```
* BD-10808.43 # show upm timers
Current Time: 2006-10-16 14:03:55
-----
UPM Profile Flags Next Execution
Timer Name time
-----
t1 p1 ep 2006-10-16 14:04:00(Every 10 secs)
timerA
-----
Flags: e - Profile is enabled, d: Profile is disabled
o -Timer is non-periodic, p - Timer is periodic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

unconfigure upm event

```
unconfigure upm event <upm-event> profile <profile-name> ports <port_list>"
```

Description

Unconfigures the event from the specified profile and port list.

Syntax Description

<upm-event>	Specifies the type of event to be unconfigured.
profile <profile-name>	Specifies the profile from which the event will be unconfigured.
ports <port-list>	Unconfigures the UPM profile from the specified port list.

Default

N/A.

Usage Guidelines

This command removes an event from the specified profile and port list.

History

This command was first available in ExtremeXOS 11.6.

unconfigure upm timer

```
unconfigure upm timer <timerName> profile <profileName>
```

Description

Removes a UPM profile from a UPM timer.

Syntax Description

<timerName>	Unconfigures the specified UPM timer and deactivates any running timer.
<profileName>	Removes the specified profile from the UPM timer.

Default

N/A.

Usage Guidelines

Use this command to unconfigure a timer setting. This command does not delete the timer.



NOTE

The specified timer will be stopped by this command, even if it has been activated.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

7 CLI Scripting Commands

This chapter describes commands for:

- CLI Scripting

By using the CLI-based scripting capability in ExtremeXOS, users can significantly automate switch management.

Users can use CLI scripting with Universal Port management to create dynamic profiles that are activated by a trigger event, and cause dynamic changes to the configuration of the switch to enforce a policy.

configure cli mode

```
configure cli mode [persistent | non-persistent]
```

Description

Configures the persistent nature of command execution.

Syntax Description

persistent	Configures command execution to be persistent.
non-persistent	Configures command execution to be not persistent.

Default

The default mode is non-persistent.

Usage Guidelines

Both persistent and non-persistent modes allow you to execute a profile. The CLI mode is implicitly set to the non-persistent mode for dynamic profiles.

Example

The following example sets command execution to be persistent:

```
configure cli mode persistent
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure cli mode scripting

```
configure cli mode scripting [abort-on-error | ignore-error]
```

Description

Configures the error handling process for CLI scripting on the switch.

Syntax Description

abort-on-error	Configures Cli scripts to be aborted if a CLI error occurs.
ignore-error	Configures the script to be executed when CLI errors occur.

Default

CLI: ignore-error

Static profiles: abort-on-error

Dynamic profiles: abort-on-error

Usage Guidelines

You can change the error-handling options within the scripts.

Example

The following command configures the switch to ignore syntax errors in CLI scripts:

```
configure cli mode scripting ignore-error
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete var key

```
delete var key <key>
```

Description

Deletes the variables that have been saved using a key.

Syntax Description

<key>	Specifies that variables associated with the key <key> must be deleted.

Default

N/A

Usage Guidelines

CLI scripting must be enabled to use this command. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following command deletes all variables associated with the key "red:"

```
delete var key red
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable cli scripting

```
disable cli scripting {permanent}
```

Description

Disables the use of CLI scripting. When used without the `permanent` option, it disables CLI scripting for the current session and is a per session setting. The `permanent` option affects new sessions only and is saved across switch reboots.

Syntax Description

<code>permanent</code>	Disables CLI scripting for new sessions only; this setting is saved across switch reboots.

Default

CLI scripting is disabled by default.

Usage Guidelines

You can disable CLI scripting for the session after it has been enabled.

Example

The following command disables CLI scripting for the current session:

```
disable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

ELSE

ELSE

Description

Command block to be executed if the condition specified in the associated `IF` statement is not met.

Syntax Description

<statements>	Actions to be executed when the conditions specified in the associated <code>IF</code> statement are not met.

Default

N/A

Usage Guidelines

CLI scripting must be enabled before using this command.

This command must be preceded by `IF <_expression> THEN <statements>` and followed by `ENDIF`.

You can insert comments by using a number sign (#).

Example

The following example will execute the `show switch` command if the value of the variable `x` is greater than 2, and execute the `show vlan` command otherwise:

```
IF ($x > 2) THEN
    show switch
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable cli scripting

```
enable cli scripting {permanent}
```

Description

Enables the use of CLI scripting. When used without the `permanent` option, it enables CLI scripting for the current session and is a per session setting. The `permanent` option enables CLI scripting for new sessions only and is saved across switch reboots.

Syntax Description

<code>permanent</code>	Enables CLI scripting for new sessions only; this setting is saved across switch reboots.

Default

CLI scripting is disabled by default.

Usage Guidelines

You must enable CLI scripting on the switch to use the scripting keywords in the script, and before you can configure or execute a script.

Example

The following command enables CLI scripting for the current session:

```
enable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

ENDIF

ENDIF

Description

Causes the IF construct to be terminated.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

The ENDIF command should be used after the IF <_expression> THEN <statement(s)> command.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

The following example will execute the `show switch` command if the value of the variable is greater than 2 and execute the `show vlan` command otherwise:

```
IF ($x > 2) THEN  
  
    show switch  
  
ELSE  
  
    show vlan  
  
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

ENDWHILE

ENDWHILE

Description

Causes the WHILE construct to be terminated.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

The `ENDWHILE` command must be used after a corresponding `WHILE <_expression> DO <statement(s)> command`.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

This example creates 10 VLANs, named x1 to x10:

```
setvar x 1

WHILE ($x <= 10) DO

    create vlan v$x

    set var x ($x + 1)

ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

IF ... THEN

```
IF (<_expression>) THEN
```

Description

Optionally executes a code block based on the condition supplied.

Syntax Description

<expression>	Specifies the condition for which the statements should be executed.
<statements>	Actions to be executed when the specified conditions are met.

Default

N/A

Usage Guidelines

This command is usually followed by statements that will be executed if the condition evaluates to true.

It may also be followed by an `ELSE` block, which will be executed if the condition evaluates to false.

The `IF` construct should be terminated by an `ENDIF` command.

The `_expression` must be enclosed in parentheses.

The `IF` construct can be nested inside other `IF` and `WHILE` constructs. Nesting is supported up to five levels. If there is incorrect nesting of `IF` conditions, an error message will be displayed. If a user tries to execute more than five nested `IF` conditions, an error message will be displayed.

The operators mentioned in [Table 23](#) can be used in an `_expression` in an `IF` condition.

You can insert comments by using a number sign (`#`).

Example

The following example will execute the `show switch` command if the value of the variable is greater than 2 and execute the `show vlan` command otherwise:

```
IF ($x > 2) THEN

    show switch

ELSE

    show vlan

ENDIF
```


History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

load var key

```
LOAD VAR KEY <key> [<var1> <var2> ...]
```

Description

Imports the specified set of variables associated with a key into the current session.

Syntax Description

<key>	Specifies the key associated with the variables to be imported.
<var1> <var1>	Specifies the variables to be imported. The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A

Usage Guidelines

The specified key should have created by the user. Also, the variables specified should have been saved using that key.

Attempting to use this command with a non-existent key will result in an error message being displayed.

Example

The following example imports the variables “username,” “ipaddr,” and “vlan” from the key “blue:”

```
load var key blue username ipaddr vlan
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

save var key

```
SAVE VAR KEY <key> [<var1> <var2> ...]
```

Description

Saves the specified variables to the specified key.

Syntax Description

<key>	Specifies the key to which the specified variables will be saved .
<var1> <var1>	Specifies the variables to save, The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A

Usage Guidelines

The variables saved by the `SAVE VAR` command are represented by the specified key and can be retrieved and restored in the context in which this profile was applied. They are available to rollback events like user-unauthenticate and device-undetected. The key option allows the user to save data for a unique key and retrieve the saved data based on this key. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following example saves the variables “username,” “ipaddr,” and “vlan” to the key “blue:”

```
save var key blue username ipaddr vlan
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

set var

```
set var <varname> <_expression>
```

Description

Creates and sets the CLI scripting variable to the desired value.

Syntax Description

<varname>	Specifies the name of the CLI scripting variable. Valid format is \$VARNAME (case insensitive, character string up to 32 characters).
<_expression>	Specifies the _expression whose value should be evaluated and used to set the variable.

Default

N/A.

Usage Guidelines

The format of a local variable (case insensitive) is: \$VARNAME

An error message will be displayed if the user attempts to use a variable name with a length greater than 32 characters.

If a variable already exists, it will be overwritten. No error message will be displayed.

Example

The following examples show some ways you can manipulate variables:

```
Set var x 100
Set var x ($x + 2)
Set var y ($x - 100)
Set var y ($(x) - 100)
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show var

```
show var {<varname>}
```

Description

Displays the current session variables or the named variable.

Syntax Description

<varname>	Displays the variable specified, if present.

Default

N/A.

Usage Guidelines

Use this command to see the list of current session variables. The display will include the variable name and value.

Example

The output of this command is similar to the following:

```
# show var
-----
Count : 6
-----

-----
variableName          variableValue
-----
CLI.SESSION_TYPE      serial
CLI.USER              admin
STATUS                0
x                     55
y                     100
-----

# show var x
-----
Name                  Value
-----
x                     55
-----
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

WHILE ... DO

```
WHILE (<_expression>) DO
```

Description

Executes a code block while the specified condition is true.

Syntax Description

<_expression>	Specifies the condition for which the statements should be executed while the condition is true.
<statements>	Set of statements to be executed while the condition is true.

Default

N/A

Usage Guidelines

This command is usually followed by statements to be executed while the condition is true and the entire construct is terminated by an `ENDWHILE` command.

The `_expression` must be enclosed in parentheses.

Nesting is supported up to five levels. An error message will be displayed if there is incorrect nesting of `WHILE` conditions. An error message will be displayed if a user tries to execute more than five `WHILE` conditions.

Ctrl-C can be used to break out of a `WHILE` loop(s). Breaking out of any number of `WHILE` loops always clears all the `WHILE` loops .

The operators mentioned in [Table 23](#) can be used in an `_expression` in a `WHILE` condition.

You can insert comments by using a number sign (#).

Example

This example creates 10 VLANs, named x1 to x10:

```
setvar x 1

WHILE ($x <= 10) DO

    create vlan v$x

    set var x ($x + 1)

ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

Connectivity Fault Management (CFM), discussed in the emerging IEEE 802.1ag specification, allows you to detect, verify, and isolate connectivity failures in virtual bridged LANs. Part of this specification is a toolset to manually check connectivity, which is sometimes referred to as Layer 2 ping.

**NOTE**

The ExtremeXOS implementation of CFM is based on draft 4.1 of the IEEE 802.1ag standard.

You create hierarchical networks, or domains, and test connectivity within that domain by sending Layer 2 messages, known as Connectivity Check Messages (CCMs). You use these domains to send loopback messages and link trace messages.

**NOTE**

CFM is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain add association integer

```
configure cfm domain <domain_name> add association integer <int> vlan  
<vlan_name>
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the 2-octet integer MA format.

Syntax Description

domain_name	Specifies the domain you want to associate with this MA.
int	Enter an integer to name the MA. The range is between 0 and 65535.
vlan_name	Specifies the VLAN you want to assign to this MA. Each VLAN belongs to only one MA, and each MA contains only one VLAN.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA. Each VLAN can be associated with a single association, or MA.

You add the MA, or association, to the domain, and the MA uses the MA level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates a 2-octet integer MA (350) that associates the domain brazil and the VLAN admin:

```
configure cfm domain brazil add association integer 350 vlan admin
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain add association string

```
configure cfm domain <domain_name> add association string <name> vlan  
<vlan_name>
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the character string MA format.

Syntax Description

domain_name	Specifies the domain you want to associate with this MA.
string	Enter up to 100 alphanumeric characters to name the MA.
vlan_name	Specifies the VLAN you want to assign to this MA. Each VLAN belongs to only one MA, and each MA contains only one VLAN.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA. Each VLAN can be associated with a single association, or MA.

You add the MA, or association, to the domain, and the MA uses the MA level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA named service that associates the MD spain and the VLAN finance:

```
configure cfm domain service add association string spain vlan finance
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain add association vpn-id oui index

```
configure cfm domain <domain_name> add association vpn-id oui <oui> index
<index> vlan <vlan_name>
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the RFC 2685 VPN ID MA format.

Syntax Description

domain_name	Specifies the domain you want to associate with this MA.
oui	Enter a virtual private network (VPN) Organizational Unique Identifier (OUI) in the format XX:XX:XX as part of the name for the MA.
index	Enter the 32-bit VPN index you want to append to the OUI to name the MA. The range is 0 to 4294967295.
vlan_name	Specifies the VLAN you want to assign to this MA. Each VLAN belongs to only one MA, and each MA contains only one VLAN.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA. Each VLAN can be associated with a single association, or MA.

You add the MA, or association, to the domain, and the MA uses the MA level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA with the VPN ID of 11:22:33 50 that associates the domain spain and the VLAN accounting:

```
configure cfm domain spain add association vpn-id oui 11:22:33 index 50 vlan
accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association add end-point

```
configure cfm domain <domain_name> association <association_name> ports
<port_list> add end-point [inward-facing | outward-facing ] <mepid>
```

Description

Configures the maintenance end points (MEPs) for that MA; use this command to configure both the inward-facing and outward-facing MEPs. Each MEP within a single MA must have a unique MEP ID.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are configuring an MEP.
ports	Enter the port number you want to configure as either in inward-facing or an outward-facing MEP.
inward-facing	Enter the port to be the inward-facing port of the MA; this MEP sends CCM messages to all ports—other than the sending switch port—in this MA on this switch.
outward-facing	Enter the port to be the outward-facing port of the MA; this MEP sends CCM messages out of the configured physical port.
mepid	Enter a value for this MEP. The range is 1 to 8191. NOTE: On each MA, each MEPID must be unique.

Default

N/A.

Usage Guidelines

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.



NOTE

Ensure that you assigned the port number correctly to the inward-facing MEP and to the outward-facing MEP, or the CCM messages go in the wrong direction.

These ports must already be in the MA (VLAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

The following port(s) <portlist> are not part of the associations VLAN.

Each MEP must have unique MEPID per MA.



NOTE

If you configured MAC-in-MAC on a port or have an untagged vMAN on that port, you cannot configure CFM on that port.

Not all combinations of MPS are allowed on the same port within an MA; only the following combinations can be on the same port within an MA:

- inward-facing MEP and MIP
- inward-facing MEP and outward-facing MEP
- outward-facing MEP and MIP

You can configure a total of 32 inward-facing MEPs and 32 outward-facing MEPs on a single switch.

Use the `show cfm` command to verify your configuration.

Example

The following command configures port 5:10 to be the inward-facing MEP on the test association in the brazil domain, with an mepid of 500:

```
configure cfm domain brazil association test ports 5:10 add end-point inward-facing 500
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association add filter-function

```
configure cfm domain <domain_name> association <association_name> ports
<port_list> add filter-function
```

Description

Configures the ports that function as CFM filters for that MA; these are referred to as CFFs. CFFs pass only those CFMs with a value superior to the MA level set for the specified MA; CFFs drop the rest of the CFMs.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are configuring a CFF.
ports	Enter the port number you want to configure as a CFF.

Default

N/A.

Usage Guidelines

These ports must already be in the MA (VLAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

The following port(s) <portlist> are not part of the associations VLAN.



NOTE

If you configured MAC-in-MAC on a port or have an untagged vMAN on that port, you cannot configure CFM on that port.

Once you assign a port as a CFF, you cannot assign it to any other MP. An unlimited number of CFFs are supported per switch.

Use the [show cfm](#) command to verify your configuration.

Example

The following command configures port 2:4 as a CFF on the 350 association in the finance domain:

```
configure cfm domain finance association 350 ports 2:4 add filter-function
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association add intermediate-point

```
configure cfm domain <domain_name> association <association_name> ports  
<port_list> add intermediate-point
```

Description

Configures the maintenance intermediate points (MIPs) for that MA. MIPs do not send CCMs nor collect or store CCM information; MIPs are used as a hop when a Layer 2 traceroute frame passes through the port.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are configuring a MIP.
ports	Enter the port number you want to configure as a MIP.

Default

N/A.

Usage Guidelines

These ports must already be in the MA (VLAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

The following port(s) <portlist> are not part of the associations VLAN.



NOTE

If you are running MAC-in-MAC on that port or an untagged vMAN port, you cannot use that port as an MP for 802.1ag.

Not all combinations of MPS are allowed on the same port within an MA; only the following combinations can be on the same port within an MA:

- inward-facing MEP and MIP
- inward-facing MEP and outward-facing MEP
- outward-facing MEP and MIP

You can configure a total of 32 MIPs on a single switch.

Use the [show cfm](#) command to verify your configuration.

Example

The following command configures port 1:20 as a MIP on the 350 association in the spain domain:

```
configure cfm domain spain association 350 ports 1:20 add intermediate-point
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association delete end-point

```
configure cfm domain <domain_name> association <association_name> ports
<port_list> delete end-point [inward-facing | outward-facing ]
```

Description

Deletes a maintenance end point (MEP) from that MA; use this command to delete either the inward-facing or the outward-facing MEP.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are deleting an MEP.
ports	Enter the port number you want to delete as either in inward-facing or an outward-facing MEP.
inward-facing	Enter the port of the inward-facing MEP you are deleting.
outward-facing	Enter the port of the outward-facing MEP you are deleting.

Default

N/A.

Usage Guidelines

Use this command to delete an MEP.

Use the [show cfm](#) command to verify your configuration.

Example

The following command deletes port 5:10 as the inward-facing MEP on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:10 delete end-point inward-facing
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association delete filter-function

```
configure cfm domain <domain_name> association <association_name> ports  
<port_list> delete filter-function
```

Description

Deletes a CFM filter function (CFF) from that MA.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are deleting a CFF.
ports	Enter the port number you want to delete as a CFF.

Default

N/A.

Usage Guidelines

Use this command to delete a CFF.

Use the [show cfm](#) command to verify your configuration.

Example

The following command deletes port 1:40 as a CFF on the test association in the spain domain:

```
configure cfm domain spain association test ports 1:40 delete filter-function
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association delete intermediate-point

```
configure cfm domain <domain_name> association <association_name> ports  
<port_list> delete intermediate-point
```

Description

Deletes a maintenance intermediate point (MIP) from that MA.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are deleting an MIP.
ports	Enter the port number you want to delete as an MIP.

Default

N/A.

Usage Guidelines

Use this command to delete an MIP.

Use the [show cfm](#) command to verify your configuration.

Example

The following command deletes port 5:12 as an MIP on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:12 delete intermediate-point
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association end-point mepid

```
configure cfm domain <domain-name> association <association_name> ports
<port_list> end-point [inward-facing | outward-facing] mepid <mepid>
```

Description

Allows you to change the MEP ID for a previously configured MEP. Each MEP within a single MA must have a unique MEP ID.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are changing an MEP ID.
ports	Enter the port number you want to change the MEP ID.
inward-facing	Enter this variable if you are changing the MEP ID on an inward-facing MEP.
outward-facing	Enter this variable if you are changing the MEP ID on an outward-facing MEP.
mepid	Enter the new value for this MEP. The range is 1 to 8191.
	NOTE: On each MA, each MEPID must be unique.

Default

N/A.

Usage Guidelines

Use this command to change the MEPID on a previously configured inward-facing or outward-facing MEP. If you attempt to change the MEPID on a port that is either not an MEP or the incorrectly facing MEP, the system returns an error message.

Use the [show cfm](#) command to verify your configuration.

Example

The following command changes the MEP ID to 75 on the previously configured port 2:4 inward-facing MEP on the 350 association in the finance domain:

```
configure cfm domain finance association 350 ports 2:4 end-point inward-facing mepid
75
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain association end-point transmit-interval

```
configure cfm domain <domain_name> association <association_name> ports
<port_list> end-point [inward-facing | outward-facing] transmit-interval
<interval>
```

Description

Allows you to change time interval for an MEP to send out a CCM. Extreme Networks recommends configuring this value as at least 1 second. The system rounds down the transmit interval to the nearest 1/100th of a second; that is, 1236 is rounded to 1230 ms.

Syntax Description

domain_name	Enter the domain associated with the MA you are configuring.
association_name	Enter the name of the MA for which you are changing the time interval an MEP sends out CCM.
ports	Enter the port number of the MEP on which you are changing the time interval it sends out a CCM.
inward-facing	Enter this variable if you are changing the time interval for sending a CCM on an inward-facing MEP.
outward-facing	Enter this variable if you are changing the time interval for sending a CCM on an outward-facing MEP.
interval	Enter the new value (in ms) you want for the transmit interval that the MEP uses to send CCMs.

Default

1000 ms.

Usage Guidelines

Use this command to change the time interval between sending out CCMs on a previously configured inward-facing or outward-facing MEP. If you attempt to change the interval on a port that is either not an MEP or the incorrectly facing MEP, the system returns a message.



NOTE

Extreme Networks recommends that you use a transmit interval of at least 1 second (1000 ms).

The system also uses this value multiplied by 3.5 as the TTL value sent with the CCMs. The receiving switch uses this value to determine when the MEP is no longer alive.

Use the `show cfm` command to verify your configuration and the `show cfm detail` command to display the configured lifetime.

Example

The following command changes the interval the inward-facing MEP (previously configured on port 2:4) uses to send CCM messages on the 350 association in the finance domain to 2 seconds:

```
configure cfm domain finance association 350 ports 2:4 end-point inward-facing  
transmit-interval 2000
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain delete association

```
configure cfm domain <domain_name> delete association <association_name>
```

Description

Deletes a maintenance association (MA), including all its configured values, from the switch.

Syntax Description

domain_name	Enter the domain associated with the MA you are deleting.
association_name	Enter the name of the MA you are deleting.

Default

N/A.

Usage Guidelines

When you delete an association, or MA, you also remove all its configured values from the switch. These values include all configured MEPs, MIPs, and CFFs.

Example

The following command deletes the MA test, in the domain of brazil, from the switch, along with all its configured MIPs, MEPs, and CFFs:

```
configure cfm domain brazil delete association test
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure cfm domain ma-level

```
configure cfm domain <domain_name> ma-level <level>
```

Description

Changes a previously configured MA level for the specified domain.

Syntax Description

domain_name	Enter the name of the domain for which you want to change the MA level.
level	Specifies the new MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level. Thus, a given MA level exists only once on a switch.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command changes the MA level of a previously created domain extreme to 2:

```
configure cfm domain extreme ma-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create cfm domain dns ma-level

```
create cfm domain dns <name> ma-level <level>
```

Description

Creates a maintenance domain (MD) in the DNS name format and assigns an MA level to that domain.

Syntax Description

name	Assigns the name you want for this domain, using the DNS name format. Enter alphanumeric characters for this format; the maximum is 100 characters.
level	Specifies the MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level.

You assign each domain a maintenance association (MA) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the lowest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the DNS name format, named extreme and assigns that domain an MA level of 2:

```
create cfm domain dns extreme ma-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create cfm domain iso3166 ma-level

```
create cfm domain iso3166 <name> ma-level <level>
```

Description

Creates a maintenance domain (MD) in the ISO 3166 country code format and assigns an MA level to that domain.

Syntax Description

name	Assigns the name you want for this domain, using the ISO 3166 country code format. Enter alphanumeric characters for this format; the maximum is 100 characters.
level	Specifies the MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level.

You assign each domain a maintenance association (MA) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the lowest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the ISO 3166 country code format, named brazil and assigns that domain an MA level of 1:

```
create cfm domain iso3166 brazil ma-level 1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create cfm domain m1400 ma-level

```
create cfm domain m1400 <name> ma-level <level>
```

Description

Creates a maintenance domain (MD) in the ITU-T M.1400 carrier code format and assigns an MA level to that domain.

Syntax Description

name	Assigns the name you want for this domain, using the ITU-T M.1400 carrier code format. Enter alphanumeric characters for this format; the maximum is 100 characters.
level	Specifies the MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level.

You assign each domain a maintenance association (MA) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the lowest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the ITU-T M.1400 carrier code format, named atlanta and assigns that domain an MA level of 7:

```
create cfm domain m1400 atlanta ma-level 7
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create cfm domain mac ma-level

```
create cfm domain mac <mac> <int> ma-level <level>
```

Description

Creates a maintenance domain (MD) in the MAC address + 2-octet integer format and assigns an MA level to that domain.

Syntax Description

mac	Enter a MAC address in the format XX:XX:XX:XX:XX:XX to specify part of the domain name.
int	Enter the 2-octet integer you want to append to the MAC address to specify the domain name.
level	Specifies the MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level.

You assign each domain a maintenance association (MA) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the lowest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the MAC + 2-octet integer format, with the MAC address of 11:22:33:44:55:66 and an integer value of 63; it also assigns that domain an MA level of 2:

```
create cfm domain mac 11:22:33:44:55:66 63 ma-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create cfm domain vpn-id oui index ma-level

```
create cfm domain vpn-id oui <oui> index <index> ma-level <level>
```

Description

Creates a maintenance domain (MD) in the RFC 2685 VPN ID format and assigns an MA level to that domain. The RFC 2685 VPN ID format consists of a 24-bit VPN OUI + a 32-bit VPN index value.

Syntax Description

oui	Enter a virtual private network (VPN) Organizational Unique Identifier (OUI) in the format XX:XX:XX to specify part of the domain name.
index	Enter the 32-bit VPN index you want to append to the OUI to specify the domain name.
level	Specifies the MA level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MA level.

You assign each domain a maintenance association (MA) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the lowest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 0 to 2 for end users
- 3 and 4 for Internet service providers (ISPs)
- 5 to 7 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the RFC 2658 VPN ID format, with a VPN OUI of 11:22:33 and a VPN index value of 50; it also assigns that domain an MA level of 1:

```
create cfm domain vpn-id oui 11:22:33 index 50 ma-level 1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

delete cfm domain

```
delete cfm domain <domain>
```

Description

Deletes the specified maintenance domain (MD) from the switch, as well as all configuration setting related to this MD.

Syntax Description

domain	Enter the name of the domain you want to delete.

Default

N/A.

Usage Guidelines

This command deletes all configuration settings related to the domain—for example, all MAs, MIPs, MEPs, and CFFs—as well as the domain itself.

Example

The following command deletes the domain atlanta (as well as all settings related to this domain):

```
delete cfm domain atlanta
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

ping mac port

```
ping mac <mac> port <port> {domain} <domain_name> {association}
<association_name>
```

Description

Allows you to ping on the Layer 2 level throughout the specified domain and MA. The ping, or loopback message (LBM), goes from the MEP configured on the port toward the given MAC address.

Syntax Description

mac	Enter the unique system MAC address on the device you want to reach. Enter this value in the format XX:XX:XX:XX:XX:XX.
port	Enter the port number of the MEP from which you are issuing the ping.
domain	Enter this keyword.
domain_name	Enter the name of the domain from which you are issuing the ping.
association	Enter this keyword.
association_name	Enter the name of the association from which you are issuing the ping.

Default

N/A.

Usage Guidelines

You must have CFM parameters configured prior to issuing a Layer 2 ping.

In order to send a Layer 2 ping, you must specify the port (MEP), the domain, and the MA from which you are issuing the ping. An inward-facing MEP sends the ping to all ports (except the sending port) on the VLAN that is assigned to the specified MA, and an outward-facing MEP sends the ping out from that port from that MA toward the specified MAC address.

All MIPs along the way forward the LBM to the destination. The destination MP responds back to the originator with a loopback reply (LBR).

This command sends out a ping from the MEP configured on the specified port toward the specified MAC address. If you attempt to send a ping message from a port that is not configured as a MEP, the system returns an error message. If the specified port is configured as both an inward-facing and an outward-facing MEP, the Layer 2 forwarding table decides which to use. If the specified MAC address is not present in the Layer 2 forwarding table (FDB), the system cannot send the ping.

Example

The following command sends a Layer 2 ping to the unique system MAC address 00:04:96:1F:A4:31 from the previously configured inward-facing MEP (port 2:4) in the speed association in the atlanta domain:

```
ping mac 00:04:96:1F:A4:31 port 2:4 atlanta speed
```

The following is sample output from the Layer 2 ping command:

```
BD-10808.24 # ping mac 00:30:48:42:F1:68 port 1:1 "dnsname" "short ma name"
Send L2 Ping inwards from 1:1, waiting for responses [press Ctrl-C to abort].
49 bytes from 00:30:48:42:f1:68, seq=10 time=9 ms
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show cfm

```
show cfm {<domain_name> {<association_name> {{ports <port_list>} {[end-
point [inward-facing | outward-facing] | intermediate-point | filter-
function]]}}}}
```

Description

Displays the current CFM configuration on the switch.

Syntax Description

domain_name	Enter the name of the domain you want to display.
association_name	Enter the name of the association you want to display.
port_list	Enter the ports in the domain and association you want to display.
inward-facing	Enter this to display the inward-facing MEP for the specified MA.
outward-facing	Enter this to display the outward-facing MEP for the specified MA.
intermediate-point	Enter this to display the MIPs for the specified MA.
filter-function	Enter this to display the CFFs for the specified MA.

Default

N/A.

Usage Guidelines

This command displays the following information:

- Domain names
- MA levels
- Association names
- VLAN names
- Inward-facing MEPs
- MEPIDs
- MEP transmit intervals
- Outward-facing MEPs
- Intermediate points (MIPs)
- Total number of CFM ports on the switch

See [“Supported Instances for CFM”](#) for the number of domains, ports, MEPs, MIPs, CFFs, and associations supported on the switch.

Example

The following command displays the current CFM configuration on the switch:

```
show cfm
```

The following is sample output from this command:

```
BD-10808.5 # show cfm
Domain: "m1400name", MA Level: 0
Domain: "isoname", MA Level: 1
Domain: "dnsname", MA Level: 2
  Association: "short ma name", VLAN "one" with 4 cfm ports
    port 1:1; Inward-facing End Point, mepid: 12, transmit-interval: 1000 ms
    port 1:2; Intermediate Point
    port 2:2; Intermediate Point
Domain: "00:11:22:33:44:55.12", MA Level: 3

Switch contains a total of 4 cfm ports.
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show cfm detail

```
show cfm {<domain_name> {<association_name> {ports <port_list>} {[end-point  
[inward-facing | outward-facing]]}}}} detail
```

Description

Displays the MEP CCM database.

Syntax Description

domain_name	Enter the name of the domain for which you want to display the MEP CCM databases.
association_name	Enter the name of the association for which you want to display the MEP CCM databases.
port_list	Enter the ports in the domain/association for which you want to display the CCM databases.
inward-facing	Enter this to display the CCM database on the inward-facing MEP for the specified MA.
outward-facing	Enter this to display the CCM database on the outward-facing MEP for the specified MA.

Default

N/A.

Usage Guidelines

If you don't specify any parameters or variables, the system displays information on all CCM databases on the switch.

This command displays the following items of the CCM database:

- The name of the domain and association
- Port number
- MP and type
- MAC address of remote end points
- MEP IDs
- Lifetime for CCM messages from each remote end point
- Actual age of CCM messages



NOTE

The TTL for the CCM messages from the MP you are working on is 3.5 times the transmission interval.

Example

The following command displays the CCM databases on the switch:

```
show cfm detail
```

The following is sample output from this command:

Domain/ Association	Port	MP	Remote End-Point MAC Address	MEP ID	Life time	Age
=====						
dnsname						
short ma name	1:1	IE	00:30:48:42:f3:12	43	3500	770
	1:1	IE	00:30:48:42:f1:68	69	3500	380
	2:1	IE	00:30:48:42:f3:12	12	3500	570
	2:1	IE	00:30:48:42:f1:68	69	3500	380
=====						

Maintenance Point: (IE) Inward-Facing End-Point, (OE) Outward-Facing End-Point
 NOTE: The Domain and Association names are truncated to 28 characters, Lifetime
 and Age are in milliseconds.

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show cfm statistics

```
show cfm {<domain_name> {<association_name> {ports <port_list>} {end-point
[inward-facing | outward-facing] | filter}}}} statistics
```

Description

Displays the MP statistics.

Syntax Description

domain_name	Enter the name of the domain for which you want to display the MP statistics.
association_name	Enter the name of the association for which you want to display the MP statistics.
port_list	Enter the ports in the domain/association for which you want to display the MP statistics.
inward-facing	Enter this to display the statistics for the inward-facing MEP for the specified MA.
outward-facing	Enter this to display the statistics for the outward-facing MEP for the specified MA.
filter	Enter this to display the statistics for the CFFs for the specified MA.

Default

N/A.

Usage Guidelines

The system displays the software counters; that is, packets blocked by hardware are not included.

To clear these counters, use the `clear counters` command.

This command displays the following information for the statistics:

- The name of the domain and association
- Port number
- MP and type
- Number of packets blocked because the MA level in the packet is *superior* to the MA level configured on the specified port (A CCM is forwarded when the embedded MA level is numerically equal to or *lower* than the configured MA level on the switch; that is, a domain configured with an MA level of 5 forwards CCMS with levels 0, 1, 2, 3,4, and 5; CCMs with a level of 6 or 7 are blocked.)
- Number of packets received with an unknown CFM message type

Example

The following command displays the CCM databases on the switch:

```
show cfm statistics
```

The following is sample output from this command:

```

Domain/      Port    MP  eff_level_in/ iff_level_in/ iff_opcode
Association  eff_level_out iff_level_out
=====
dnsname
short ma name      2:1    IE  0           0           0
                   0
                   2:1    OE  0           0           0
                   2:2    FF 396
                   0
=====
Maintenance Point: (IE) Inward-Facing End-Point, (I) Intermediate-Point,
                   (OE) Outward-Facing End-Point, (FF) Filter-Function
NOTE: The Domain and Association names are truncated to 28 characters, also the
      counters displayed are software counters, i.e. packets blocked by
      hardware are not included.
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

traceroute mac port

```
traceroute mac <mac> {inward-facing-end-point} port <port> {domain}
<domain_name> {association} <association_name> {ttl <ttl>}
```

Description

Allows you to send out a Link Trace Message (LTM) for the specified MA from the MEP configured on the port for the specified MAC address to the end of the MA.

Syntax Description

mac	Enter the unique system MAC address on the port configured as a MEP for the specified MA. Enter this value in the format XX:XX:XX:XX:XX:XX.
inward-facing-end-point	Use this keyword to force the LTM to be send from an inward-facing MEP if both an outward-facing MEP and an inward-facing MEP are configured on the same port.
port	Enter the port number of the MEP from which you are issuing the LTM.
domain	Enter this keyword.
domain_name	Enter the name of the domain from which you are issuing the ping.
association	Enter this keyword.
association_name	Enter the name of the association from which you are issuing the ping.
ttl	Enter this keyword.
ttl	Enter the upper limit of MIPs the LTM can pass prior to reaching its destination.

Default

TTL default value is 64.

Usage Guidelines

Use this command to sent an LTM from the MEP on the port for the given MAC address. If no MEP is configured on the port, the system returns an error message.

If both an inward-facing and outward-facing MEP are configured on the same port, the system uses the outward-facing MEP. If you want to use the inward-facing MEP in this situation, enter the `inward-facing-end-point` keyword. After you issue the command, the system prints out the route the LTM message took.

Each MIP along the route passes the LTM along only in the direction of the path and sends a packet back to the originating MAC notifying that it passed the LTM.

This command returns the following information for each point along the route:

- TTL value
- CFM source MAC address
- Reply
- Reply ingress—inward-facing MEP
- Reply egress—outward-facing MEP

Example

The following command sends an LTM:

```
traceroute MAC 11:22:33:44:55:66 port 2:4 domain atlanta association 350
```

The following is sample output from this command:

```
TTL  CFM Source MAC      Reply  Reply Ingress      Reply Egress
=====
63   00:30:48:42:f3:12  F --f-              00:30:48:42:f3:12  -----o
62   00:30:48:42:f1:68  - nu--  00:30:48:42:f1:68  --
=====
Reply Flags : (F) FwdYes, (n) RlyNone, (u) RlyUnknown, (f) RlyFDB,
              (m) RlyMEPDB
Ingress Flags: (o) IngOk, (b) IngBlocked
Egress Flags : (n) EgrNone, (t) EgrTTL, (d) EgrDown, (b) EgrBlocked,
              (g) EgrGVRP, (o) EgrOK
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

Beginning with ExtremeXOS version 11.2, the software supports the Link Layer Discovery Protocol (LLDP). LLDP provides a standard method of discovering and representing the physical network connections of a given network management domain. LLDP works together with Extreme Discovery Protocol (EDP); it also works independently, you do not have to run EDP to use LLDP. The LLDP neighbor discovery protocol allows you to discover and maintain accurate network topologies in a multivendor environment.

The one-way protocol LLDP transmits periodic advertisements containing device information and media-specific configuration information to neighbors attached to the same network. LLDP agents cannot solicit information from other agents by way of this protocol. The type length value (TLV) format with link layer control frames is used to communicate with other LLDP agents. LLDP agents also receive link layer control frames, extract the information in TLVs, and store them in LLDP MIB objects.

You configure LLDP per port, and each port can store received information for a maximum of four neighboring ports. All LLDP configurations are saved across reboots when you issue the `save configuration` command.



NOTE

Because each LLDP protocol data unit (LLDPDU) is limited to 1500 bytes, the system drops configured optional TLVs that exceed this limit. To see if your system is dropping TLVs, issue the `show lldp statistics` command.

Beginning with ExtremeXOS version 11.5, you can configure the switch for the following additional types of LLDP TLVs:

- Avaya-Extreme Networks proprietary TLVs

These TLVs are proprietary to Avaya and Extreme Networks; they do not work on other vendors' devices. Some proprietary TLVs can only be received by the switch; they cannot be transmitted. (Refer to the *ExtremeXOS Concepts Guide* for a complete listing of these TLVs.)

- LLDP MED TLVs

These are the LLDP media endpoint discovery (MED) messages. After they are enabled, the LLDP MED TLVs are transmitted only *after* the switch *detects* a neighbor transmitting LLDP MED TLVs. Some LLDP MED TLVs can only be received by the switch; they cannot be transmitted. (Refer to the *ExtremeXOS Concepts Guide* for a complete listing of these TLVs.)

configure lldp med fast-start repeat-count

```
configure lldp med fast-start repeat-count <count>
```

Description

The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV. This command configures how many times, from 1 to 10, the switch sends out an LLDP MED packet with an interval of 1 second.

Syntax Description

count	Specifies the number of times the switch transmits LLDP MED TLVs each second (once it detects a neighbor transmitting LLDP MED TLVs). The range is 1 to 10.

Default

3.

Usage Guidelines

When the switch detects a MED-capable device, this count determines how many times the switch sends a LLDP MED TLVs with an interval of 1 second. The fast-start feature enables the MED-capable device to quickly learn information; this command changes the value from the default 3. The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV.



NOTE

After you configure the LLDP MED capability TLV, the fast-start feature automatically runs. To configure the LLDP MED capability TLV, use the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

Example

The following command configures fast learning on the switch to a value of 2:

```
configure lldp med fast-start repeat-count 2
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports management-address

```
configure lldp ports [all | <port_list>] [advertise | no-advertise]
management-address
```

Description

Configures the LLDP port to advertise or not to advertise management address information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

With ExtremeXOS, you can only add one management address TLV per LLDPDU and the information must be the IP address configured on the management VLAN. If no IP address is assigned to the management VLAN, the system sends the system MAC address. LLDP does not send out IPv6 addresses in this field.

Example

The following command advertises the management address information for port 1:5:

```
configure lldp ports 1:5 advertise management-address
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports port-description

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] port-  
description
```

Description

Configures the LLDP port to advertise or not advertise port description information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:7 to not advertise the port description information to neighbors:

```
configure lldp ports 1:7 no-advertise port-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-capabilities

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-  
capabilities
```

Description

Configures the LLDP port to advertise or not to advertise its system capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When at least one VLAN exists with more than two ports, bridging is sent to enabled.

When at least one VLAN on the switch has IP forwarding enabled, the system automatically sets the router bit.

Example

The following command configures all ports to advertise system capability information to neighbors:

```
configure lldp ports all advertise system-capabilities
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-description

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-  
description
```

Description

Configures the LLDP port to advertise or not to advertise its system description to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

Advertise.

Usage Guidelines

Although not mandatory according to the standard, this TLV is included in the LLDPDU by default when you enable LLDP.

When enabled, the system sends the following image (from the `show version` command) in the system description TLV:

```
ExtremeXOS version 11.2.0.12 v1120b12 by release-manager  
on Fri Mar 18 16:01:08 PST 2005
```

Example

The following command configures port 1:4 through port 1:8 to not advertise the system description information to neighbors:

```
configure lldp ports 1:4 - 1:8 no-advertise system-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-name

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-name
```

Description

Configures the LLDP port to advertise or not to advertise its system name to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:6 to advertise the system name to neighbors:

```
configure lldp ports 1:4 - 1:8 advertise system-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme call-server

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific avaya-extreme call-server <ip_address_1> {<ip_address_2>
{<ip_address_3> {<ip_address_4> {<ip_address_5> {<ip_address_6>
{<ip_address_7> {ip_address_8>}}}}}}}}
```

Description

Configures the LLDP port to advertise or not advertise up to 8 call server IP addresses to its neighbors. The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the call server(s) to use.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
ip_address_1....8	Specifies IP address of up to 8 call servers.
	NOTE: This parameter does not apply when you configure the no-advertise parameter.

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 8 call servers in a single TLV.

Example

The following command configures ports 1-5 to advertise two call server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server
10.10.10.10 10.11.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme dot1q-framing

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific avaya-extreme dot1q-framing [tagged | untagged | auto]
```

Description

Configures the LLDP port to advertise or not advertise the 802.1q framing configuration to its neighbors. The Avaya phone uses this proprietary LLDP TLV information. In addition to this LLDP TLV, you must enable LLLDP as well as configure both the LLDP MED capabilities TLV and the LLDP network policy TLV.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
tagged	Specifies to use tagging.
	NOTE: This parameter applies only when you use the <code>advertise</code> parameter.
untagged	Specifies <i>not</i> to use tagging.
	NOTE: This parameter applies only when you use the <code>advertise</code> parameter.
auto	Specifies following a predetermined sequence (see “Usage Guidelines” below).
	NOTE: This parameter applies only when you use the <code>advertise</code> parameter.

Default

No advertise.

Usage Guidelines

Before configuring this LLDP TLV, you must take the following steps:

- Enable LLDP using the `enable lldp ports` command.
- Enable the LLDP MED capabilities TLV using the `configure lldp ports vendor-specific med capabilities` command.
- Enable the LLDP MED network policy TLV using the `configure lldp ports vendor-specific med policy application` command.

This TLV is used to exchange information about Layer 2 priority tagging between the network connectivity device (switch) and the Avaya phone.

If you configure the TLV to advertise tagging, the phone uses tagging information, which it retrieves from the `configure lldp ports vendor-specific med policy application` command. If you configure the TLV to advertise untagged, the phone will not use any tagging, including 802.1q priority tagging.

If you configure the TLV to advertise auto, the phone cycles through the following sequence until an action is successful:

- Uses the configuration advertised by the LLDP MED network policy TLV, as configured by the `configure lldp ports vendor-specific med policy application` command.
- Uses the priority tagged frames configured by the phone's server.
- Sends the traffic untagged.

Example

The following command configures all ports to advertise the dot1q framing as untagged to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme dot1q-framing untagged
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme file-server

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific avaya-extreme file-server <ip_address_1> {<ip_address_2>
{<ip_address_3> {<ip_address_4>}}}
```

Description

Configures the LLDP port to advertise or not advertise up to 4 file server IP addresses to its neighbors. The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the file server(s) to use.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
ip_address_1...4	Specifies IP address of up to 4 file servers.
	NOTE: This parameter does not apply when you configure the no-advertise parameter.

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 4 file servers in a single TLV.

Example

The following command configures all ports to advertise two file server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server
10.20.10.10 10.12.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme poe-conservation-request

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific avaya-extreme poe-conservation-request
```

Description

Configures the LLDP port to advertise or not advertise a requested conservation level. By default, the requested conservation value on this proprietary LLDP TLV is 0, which is no power conservation. This LLDP TLV is sent out only on PoE-capable Ethernet ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The switch sends this proprietary LLDP TLV to request a PD to go into a certain power conservation level or request the PD to go to the maximum conservation level. This LLDP TLV is transmitted only on PoE-capable ports.

When configured to advertise, the switch sends this TLV with a requested conservation power level of 0, which requests no power conservation. To temporarily change this conservation level, use the SNMP `lldpXAvExLocPortXPoEPSEPortReqLevel` object to set a new value; but the reconfigured value will *not* be saved over a reboot. (This SNMP object can be set from 0 to 243 or 255.)

Example

The following command configures all ports to advertise the currently requested conservation level to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme poe-conservation-request
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 port-vlan-ID

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot1 port-vlan-ID
```

Description

Configures the LLDP port to advertise or not advertise port vlan ID information to its neighbors. This allows a VLAN bridge port to advertise the port VLAN identifier that is associated with untagged or priority-tagged frames.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The port VLAN ID TLV allows the port to transmit the VLAN ID associated with untagged VLANs. There can be only one port VLAN ID in each LLDPDU.

If no untagged VLANs are configured on the specified port, the TLV is not added to the LLDPDU, even if you configured this to advertise.

Example

The following command configures all ports to advertise port vlan ID information to neighbors:

```
configure lldp ports all advertise vendor-specific dot1 port-vlan-ID
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 port-protocol-vlan-ID

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot1 port-protocol-vlan-ID {vlan [all | <vlan_name>]}
```

Description

Configures the LLDP port to advertise or not advertise port VLAN information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
all	Specifies all VLANs on the port.
vlan_name	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

When configured to advertise, the switch inserts a port and protocol VLAN ID TLV for each VLAN configured on the ports. The port and protocol VLAN ID TLV allows the port to advertise if it supports protocol and/or tagged VLANs, along with the associated tagged values. A separate TLV is sent for each VLAN that you want to advertise.

By default, once you configure this TLV, the system sends all protocol-based VLANs on the port. However, the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the *specified* VLANs.



NOTE

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

This TLV does not send information on the type of protocol that the VLAN has enabled; it just says whether the port is enabled or disabled for protocol-based VLANs. As Extreme Networks devices are always capable of supporting protocol-based VLANs, once you configure this TLV, the system *always* advertises support these VLANs.

Example

The following command configures all ports to advertise port and protocol VLAN information to neighbors for all VLANs on all ports:

```
configure lldp ports all advertise vendor-specific dot1 port-protocol-vlan-id
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 vlan-name

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific dot1 vlan-name {vlan [all | <vlan_name>]}
```

Description

Configures the LLDP port to advertise or not advertise VLAN name information to its neighbors. Use this TLV to advertise information for the tagged VLANs you want to specify on the port. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
vlan	Specifies all VLANs on the port.
vlan_name	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

The VLAN name TLV sends the VLAN name and the tag used; it associates a name to a tag for the specified VLAN. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

You can enable this TLV for tagged and untagged VLANs. When you enable this TLV for tagged VLANs, the TLV advertises the IEEE 802.1Q tag for that VLAN. (For untagged VLANs, the internal tag is advertised.) You can specify exactly *which* VLANs to advertise.

When configured to advertise, the switch inserts a VLAN name TLV for every VLAN configured on the ports. By default, once you configure this TLV, the system sends all VLAN names on the port. However, each VLAN name can require up to 32 bytes and the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the *specified* VLANs, using the keyword `vlan_name`.



NOTE

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

Example

The following command configures all ports to not advertise VLAN name information to neighbors:

```
configure lldp ports all no-advertise vendor-specific dot1 vlan-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 link-aggregation

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3 link-aggregation
```

Description

Configures the LLDP port to advertise or not advertise link-aggregation capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, this TLV is added to each LLDP port LLDPDU indicating the link-aggregation capabilities, status, and value of the master port of the load-sharing group.

Example

The following command configures port 1:12 to not advertise link-aggregation capabilities to neighbors:

```
configure lldp ports 1:12 no-advertise vendor-specific dot3 link-aggregation
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 mac-phy

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3 mac-phy
```

Description

Configures the LLDP port to advertise or not advertise MAC and physical layer capabilities to its neighbors. The capabilities include duplex and bit rate.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system add information about the speed capabilities, as well as autonegotiation support and status, of the LLDP port.

Example

The following command configures all ports to advertise MAC/PHY capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 mac-phy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 max-frame-size

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3 max-frame-size
```

Description

Configures the LLDP port to advertise or not advertise its maximum frame size to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When jumbo frames are not enabled on the specified port, the TLV reports a value of 1518 once you configure it to advertise. If jumbo frames are enabled, the TLV inserts the configured value for the jumbo frames.

Example

The following command configures ports 1:12 and 1:13 to advertise the maximum frame size to neighbors:

```
configure lldp ports 1:12 - 1:13 advertise vendor-specific dot3 max-frame-size
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 power-via-mdi

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific dot3 power-via-mdi
```

Description

Configures the LLDP port to advertise or not advertise Power over Ethernet (PoE) capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system includes this TLV. Extreme Networks recommends enabling this TLV only on PoE-capable ports.

The following information is transmitted for LLDP ports with this TLV:

- Support PoE or not
- Port class
 - Power sourcing equipment (PSE)
 - Powered device (PD)
- Power pairs used to supply power
 - Signal
 - Spare
- Power status
- Support pairs control or not
- Power class
 - Class0
 - Class1
 - Class2
 - Class2
 - Class3
 - Class4

**NOTE**

For more information on advertising power support, see the [configure lldp ports vendor-specific med power-via-mdi](#) command.

Example

The following command configures all ports to advertise power capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 power-via-mdi
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med capabilities

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific med capabilities
```

Description

Configures the LLDP port to advertise or not advertise MED capabilities. This TLV must be enabled before any of the other MED TLVs can be enabled. Also, this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

This command enables the LLDP media endpoint discovery (MED) capabilities TLV, which allows LLDP-MED network connectivity devices to definitively determine that particular endpoints support LLDP MED, and if so, to discover which LLDP MED TLVs the particular endpoint devices are capable of supporting and to which specific device class the device belongs to.

This TLV must be enabled before any of the other MED TLVs can be enabled; and this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



NOTE

Network connectivity devices wait to detect LLDP MED TLVs from endpoints before they send out LLDP MED TLVs; so 2 network connectivity devices will not exchange LLDP MED messages.

The following information is included in the LLDP MED capabilities TLV when it is transmitted:

- The supported LLDP MED TLVs—For Extreme Networks devices, these are capabilities, network policy, location, and extended power (extended power only advertised only on PoE-capable ports).
- The MED device type—For Extreme Networks devices, this is advertised as a network connectivity device (set to 4).

Example

The following command configures all ports to advertise MED capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific med capabilities
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med location-identification

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific med location-identification [coordinate-based <hex_value> | civic-
based <hex_value> | ecs-elin <elin>]
```

Description

Configures the LLDP port to advertise or not advertise MED location information. You configure up to 3 different location identifiers.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
advertise	Specifies to send the information to neighbors.
coordinate-based	Specifies using the coordinate-based location identifier. This value is exactly 16 bytes long; see RFC 3825 for details.
hex_value	Enter a hexadecimal value with each byte separated by a colon. Or, you may obtain this value from a network management application. NOTE: This parameter is not used when the no-advertise parameter is configured.
civic-based	Specifies using the civic-based location identifier. This value must have a minimum length of 6 bytes; see RFC3825 for details.
ecs-elin	Specifies using the ecs location identifier. (Emergency Call Service, as defined in the TIA-TSB-146.)
elin	Enter a numerical string; the range is 10 to 25 characters. Or, you may obtain this value from a network management application. (See the TIA-TSB-146 standard for a definition of these numbers; also, the network management application must be able to handle the LLDP MED MIB.) NOTE: This parameter is not used when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

You may need to use a specific format for your specific VoIP implementation; see the VoIP manufacturer's manual for details.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

Example

The following command configures all ports to advertise MED location information to neighbors using the ECS format:

```
configure lldp ports all advertise vendor-specific med location-identification ecs-  
elin 423233455676
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med policy application

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific med policy application [voice | voice-signaling | guest-voice |
guest-voice-signaling | softphone-voice | video-conferencing | streaming-
video | video-signaling] vlan <vlan_name> dscp <dscp_value> {priority-
tagged}
```

Description

Configures the LLDP port to advertise or not advertise MED network policy TLVs. This TLV advertises VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific applications on that port. You can advertise up to 8 TLVs, each for a specific application, per port/VLAN. Each application type can exist only *once* per port. This TLV tells the endpoint the specific VLAN to use for the specific application, along with its unique priority.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
advertise	Specifies to send the information to neighbors.
voice	Specifies voice application on specified port/VLAN(s).
voice-signaling	Specifies voice signaling application on specified port/VLAN(s).
guest-voice	Specifies guest voice application on specified port/VLAN(s).
guest-voice-signaling	Specifies guest voice signaling application on specified port/VLAN(s).
softphone-voice	Specifies soft phone voice application on specified port/VLAN(s).
video-conferencing	Specifies videoconferencing application on specified port/VLAN(s).
streaming-video	Specifies streaming video application on specified port/VLAN(s).
video-signaling	Specifies video signaling application on specified port/VLAN(s).
vlan_name	Specifies the VLAN the specified application is using. NOTE: This parameter does not apply when the no-advertise parameter is configured.
dscp_value	Specifies the DSCP value for the specified application. This is a 6-bit value from 0 to 63. NOTE: This parameter does not apply when the no-advertise parameter is configured.
priority-tagged	Use this if you want priority tagging, and the VLAN is configured as untagged on the port. (The endpoint sends out frames for the specified application with a tag of 0.) NOTE: This parameter does not apply when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

This command enables the LLDP MED network policy TLV, which allows network connectivity devices and endpoint devices to advertise VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific application on that port. This TLV can be enabled on a per port/VLAN basis. Each application type can exist only once on a port.

You can enable the transmission of a TLV policy for each application. A maximum of 8 TLVs can be enabled, and each can have a unique DSCP value and/or priority tagging.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

The following information is transmitted for LLDP ports with this TLV:

- Application type
Used as configured.
- Unknown policy flag
Set to 0.
- Tagged flag
Set to tagged for tagged VLANs; set to untagged for untagged VLANs. By default, set to 0.
- VLAN ID
Copied from the VLAN. However, if you configure the `priority-tagged` parameter, this value is set to 0.
- Layer 2 priority
Copied from the VLAN priority.
- DSCP value
Uses the value configured in the `dscp` parameter.



NOTE

See the documentation provided by the manufacturer of connected devices regarding values.

Example

The following command configures all ports to advertise videoconferencing on the VLAN video with a DSCP of 7 to neighbors:

```
configure lldp ports all advertise vendor-specific med policy application video-
conferencing vlan video dscp 7
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med power-via-mdi

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-
specific med power-via-mdi
```

Description

Configures the LLDP port to advertise or not advertise MED power requirement details. This TLV can only be enabled on a PoE-capable port and is used for advanced power management between the MED network connectivity and endpoint devices.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When enabled, this LLDP MED TLV advertises fine-grained power requirement details about PoE settings and support. This TLV can be enabled only on a PoE-capable port; the switch returns an error message if this TLV is configured for a non-PoE-capable port.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



NOTE

For additional information on power support, see the `configure lldp ports vendor-specific dot3 power-via-mdi` command.

The following information is transmitted for LLDP MED PoE-capable ports with this TLV:

- Power type
Set to PSE.
- Power source
Set to primary power source.

- Power priority
Taken from PoE port configuration.
- Power value
Taken from PoE port configuration.

Example

The following command configures all ports to advertise MED power information to neighbors:

```
configure lldp ports all advertise vendor-specific med power-via-mdi
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp reinitialize-delay

```
configure lldp reinitialize-delay <seconds>
```

Description

Configures the delay before the receive state machine is reinstalled once the LLDP transmit mode has been disabled.

Syntax Description

seconds	Specifies the delay that applies to the reinitialization attempt. The range is 1 to 10 seconds.

Default

2 seconds.

Usage Guidelines

N/A.

Example

The following command configures a reinitialization delay of 10 seconds:

```
configure lldp reinitialize-delay 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp snmp-notification-interval

```
configure lldp snmp-notification-interval <seconds>
```

Description

Configures the allowed interval at which Simple Network Management Protocol (SNMP) notifications are sent.

Syntax Description

seconds	Specifies the interval at which LLDP SNMP notifications are sent. The range is 5 to 3600 seconds.

Default

5 seconds.

Usage Guidelines

This is a global timer. If one port sends a notification, no notifications for other ports go out for the configured interval.

Example

The following command configures an interval of 60 seconds for LLDP SNMP notifications:

```
configure lldp snmp-notification-interval 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-delay

```
configure lldp transmit-delay [ auto | <seconds>]
```

Description

Configures the delay time between successive frame transmissions initiated by a value change or status change in any of the LLDP local systems Management Information Base (MIB). The `auto` option uses a formula ($0.25 * \text{transmit-interval}$) to calculate the number of seconds.

Syntax Description

<code>auto</code>	Uses the formula ($0.25 * \text{transmit-interval}$) to calculate the seconds.
<code>seconds</code>	Specifies the interval at which LLDP notifications are sent. The range is 1 to 8291.

Default

2 seconds.

Usage Guidelines

This is the timer between triggered updates.

Example

The following command configures the delay between LLDP frame transmissions for triggered updates to be automatically calculated:

```
configure lldp transmit-delay auto
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-hold

```
configure lldp transmit-hold <hold>
```

Description

Calculates the actual time-to-live (TTL) value used in the LLDPDU messages. The formula is `transmit-interval * transmit-hold`; by default the TTL value is (30*4) 120 seconds.

Syntax Description

hold	Used to calculate the TTL value; the range is 2 to 10.

Default

4.

Usage Guidelines

N/A.

Example

The following command configures the transmit-hold value (which is used to calculate the TTL of the LLDP packets) to 5:

```
configure lldp transmit-hold 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-interval

```
configure lldp transmit-interval <seconds>
```

Description

Configures the periodic transmittal interval for LLDPDU.

Syntax Description

seconds	Specifies the time between LLDPDU transmissions. The range is 5 to 32768.

Default

30 seconds.

Usage Guidelines

N/A.

Example

The following command configures a transmittal interval of 20 seconds for LLDPDU.

```
configure lldp transmit-interval 20
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable lldp ports

```
disable lldp ports [all | <port_list>] {receive-only | transmit-only}
```

Description

Disables LLDP transmit mode, receive mode, or transmit and receive mode on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
receive-only	Specifies that only the receive mode for LLDP is disabled.
transmit-only	Specifies that only the transmit mode for LLDP is disabled.

Default

Disabled.

Usage Guidelines

If you do not specify an option, both LLDP modes (transmit and receive) are disabled.

Example

The following example disables the LLDP receive mode on ports 1:2 to 1:6.

```
disable lldp ports 1:2-1:6 receive-only
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable snmp traps lldp

```
disable snmp traps lldp {ports [all | <port_list>]}
```

Description

Disables the sending of LLDP-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP traps from all ports on the switch.

Example

The following example disables sending LLDP SNMP traps on all switch ports:

```
disable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable snmp traps lldp-med

```
disable snmp traps lldp-med {ports [all | <port_list>]}
```

Description

Disables the sending of LLDP MED-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP MED traps from all ports on the switch.

Example

The following example disables sending LLDP MED SNMP traps on all switch ports:

```
disable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

enable lldp ports

```
enable lldp ports [all | <port_list>] {receive-only | transmit-only}
```

Description

Enables LLDP transmit mode, receive mode, or transmit and receive mode. If the transmit-only or receive-only option is not specified, both transmit and receive modes are enabled.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
receive-only	Specifies that the port only receives LLDP messages.
transmit-only	Specifies that the port only transmits LLDP messages.

Default

Disabled.

Usage Guidelines

If you do not specify an option, the port is enabled to both transmit and receive LLDP messages.

Once the port is enabled for LLDP in one mode and you issue another `enable lldp ports` command for another mode, that second mode replaces the original mode. For example, you may originally enable several ports to only receive LLDP messages and then want those ports to both receive and transmit LLDP messages. In that case, you issue the `enable lldp ports` command with no variables (and the receive-and-transmit mode replaces the receive-only mode).

To verify the port setting for LLDP, use the `show lldp {port [all | <port_list>]} {detailed}` command.

Example

The following example enables LLDP transmit and receive mode on port 1:4.

```
enable lldp port 1:4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable snmp traps lldp

```
enable snmp traps lldp {ports [all | <port_list>]}
```

Description

Enables the transmission of LLDP SNMP trap notifications.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines



NOTE

To enable SNMP traps for LLDP MED TLVs, you must issue a separate command; use the `enable snmp traps lldp-med {ports [all | <port_list>]}`.

If you do not specify any ports, the system sends LLDP traps for all ports.



NOTE

The Avaya-Extreme proprietary TLVs do not send traps.

Example

The following command enables LLDP SNMP traps for all ports:

```
enable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable snmp traps lldp-med

```
enable snmp traps lldp-med {ports [all | <port_list>]}
```

Description

Enables the transmission of LLDP SNMP trap notifications related to LLDP MED extension TLVs.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system sends LLDP-MED traps for all ports.

Example

The following command enables LLDP-MED SNMP traps for all ports:

```
enable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show lldp

```
show lldp {port [all | <port_list>]} {detailed}
```

Description

Displays LLDP configuration information for the specified port or ports. Use the `detailed` keyword to display the configured VLANs on the port and the enabled VLAN-specific TLVs.

Syntax Description

<code>all</code>	Specifies all ports on the switch.
<code>port_list</code>	Specifies one or more ports or slots and ports.
<code>detailed</code>	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

Use the `detailed` variable to display information regarding configured VLANs on the ports and any enabled VLAN-specific TLVs.

Example

The following example displays LLDP configuration information for the switch:

```
# show lldp
```

```
LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4   (used TTL = 120 seconds)
LLDP transmit delay        : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4
```

LLDP Port Configuration:

Port	Rx Mode	Tx Mode	SNMP Notification	Optional enabled transmit TLVs
				LLDP 802.1 802.3 MED AvEx
1:1	Enabled	Enabled	--	--D-- --- ---- CLP- ----
1:2	Enabled	Enabled	L-	--D-- --- ---- C-P- ----
7:1	Enabled	Enabled	LM	--D-- --- ---- CLP- ----

```
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags   : (P) Port Description, (N) System Name, (D) System Description
               (C) System Capabilities, (M) Mgmt Address
```

```

802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
               (L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
               (L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server
               (Q) 802.1Q Framing

```

The following example includes detailed information on the LLDP configuration for port 1:1:

```
# show lldp port 1:1 detailed
```

```

LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4   (used TTL = 120 seconds)
LLDP transmit delay         : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4

```

LLDP Port Configuration:

Port	Rx Mode	Tx Mode	SNMP Notification	Optional enabled transmit TLVs
				LLDP 802.1 802.3 MED AvEx
1:1	Enabled	Enabled	--	--D-- --- ---- CLP- ----
VLAN: Default				
VLAN: voice				
AvEx Call-Server: IP Address(es)=10.0.0.20, 10.0.0.21				
AvEx File-Server: IP Address(es)=10.0.0.20, 10.0.0.21, 10.0.0.22				
AvEx 802.1Q Framing: Mode=tagged				
MED LCI: Location Format=ECS ELIN based				
1234567890				
MED Policy: Application=voice				
VLAN=voice, DSCP=40				

```

=====
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags   : (P) Port Description, (N) System Name, (D) System Description
               (C) System Capabilities, (M) Mgmt Address
802.1 Flags  : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags  : (M) MAC/PHY Configuration/Status, (P) Power via MDI
               (L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
               (L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server
               (Q) 802.1Q Framing

```

History

This command was first available in ExtremeXOS 11.2.

The information on fast-start repeat count, MED, AvEx, and notification was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show lldp neighbors

```
show lldp {port [all | <port_list>]} neighbors {detailed}
```

Description

Displays the information related to the LLDP neighbors detected on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
detailed	Shows detailed information on the neighbors.

Default

N/A.

Usage Guidelines

You must use the `detailed` parameter to display detailed information about the received LLDP TLVs.

Example

The following example displays LLDP neighbor information for all switch ports:

```
# show lldp port all neighbors
```

Port	Neighbor Chassis ID	Neighbor Port ID	TTL	Age
1:2	00:04:96:26:A4:70	1:1	120	7
2:6	(5.1)10.201.41.146	00:04:0D:EC:EA:5C	120	3
2:7	(5.1)10.201.41.147	00:04:0D:ED:41:9B	120	3
2:10	00:01:30:F9:9E:80	8:10	120	15

NOTE: The Chassis ID and/or Port ID might be truncated to fit the screen.

The following command lists detailed LLDP neighbor information for all switch ports:

```
# show lldp all neighbors detailed
```

```
-----
LLDP Port 1:2 detected 1 neighbor
  Neighbor: 00:04:96:26:A4:70/1:1, age 12 seconds
    - Chassis ID type: MAC address (4)
      Chassis ID      : 00:04:96:26:A4:70
    - Port ID type: ifName (5)
      Port ID        : "1:1"
    - Time To Live: 120 seconds
    - System Description: "ExtremeXOS version 12.0.0.6 v1200b6 by release-ma\
```

nager on Mon Mar 19 00:37:59 PDT 2007"

LLDP Port 2:6 detected 1 neighbor

Neighbor: (5.1)10.201.41.146/00:04:0D:EC:EA:5C, age 8 seconds

- Chassis ID type: Network address (5); Address type: IPv4 (1)
Chassis ID : 10.201.41.146
- Port ID type: MAC address (3)
Port ID : 00:04:0D:EC:EA:5C
- Time To Live: 120 seconds
- System Name: "AVAECEA5C"
- System Capabilities : "Bridge, Telephone"
Enabled Capabilities: "Bridge, Telephone"
- Management Address Subtype: IPv4 (1)
Management Address : 10.201.41.146
Interface Number Subtype : System Port Number (3)
Interface Number : 1
Object ID String : "1.3.6.1.4.1.6889.1.69.2.3"
- IEEE802.3 MAC/PHY Configuration/Status
Auto-negotiation : Supported, Enabled (0x03)
Operational MAU Type : 100BaseTXFD (16)
- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
MED Device Type : Endpoint Class III (3)
- MED Network Policy
Application Type : Voice (1)
Policy Flags : Known Policy, Tagged (0x1)
VLAN ID : 0
L2 Priority : 6
DSCP Value : 46
- MED Hardware Revision: "9650D01A"
- MED Firmware Revision: "hb96xxua1_20r30s.bin"
- MED Software Revision: "ha96xxua1_20r30s.bin"
- MED Serial Number: "06N537900335"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9650"
- Avaya/Extreme Conservation Level Support
Current Conservation Level: 0
Typical Power Value : 0.0 Watts
Maximum Power Value : 0.0 Watts
Conservation Power Level : 1=0.0W
- Avaya/Extreme Call Server(s): 69.26.36.53
- Avaya/Extreme IP Phone Address: 10.201.41.146 255.255.255.0
Default Gateway Address : 10.201.41.1
- Avaya/Extreme CNA Server: 0.0.0.0
- Avaya/Extreme File Server(s): 10.201.41.36
- Avaya/Extreme IEEE 802.1q Framing: Tagged

LLDP Port 2:7 detected 1 neighbor

Neighbor: (5.1)10.201.41.147/00:04:0D:ED:41:9B, age 8 seconds

- Chassis ID type: Network address (5); Address type: IPv4 (1)
Chassis ID : 10.201.41.147
- Port ID type: MAC address (3)
Port ID : 00:04:0D:ED:41:9B
- Time To Live: 120 seconds
- System Name: "AVAED419B"
- System Capabilities : "Telephone"

```

    Enabled Capabilities: "Telephone"
- Management Address Subtype: IPv4 (1)
  Management Address       : 10.201.41.147
  Interface Number Subtype : System Port Number (3)
  Interface Number         : 1
  Object ID String         : "1.3.6.1.4.1.6889.1.69.2.5"
- IEEE802.3 MAC/PHY Configuration/Status
  Auto-negotiation         : Supported, Enabled (0x03)
  Operational MAU Type     : 100BaseTXFD (16)
- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
  MED Device Type          : Endpoint Class III (3)
- MED Network Policy
  Application Type         : Voice (1)
  Policy Flags             : Known Policy, Tagged (0x1)
  VLAN ID                  : 0
  L2 Priority               : 6
  DSCP Value               : 46
- MED Hardware Revision: "9610D01A"
- MED Firmware Revision: "hb96xxua1_20r30s.bin"
- MED Software Revision: "ha96xxua1_20r30s.bin"
- MED Serial Number: "06N538825133"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9610"
- Avaya/Extreme Conservation Level Support
  Current Conservation Level: 0
  Typical Power Value       : 0.0 Watts
  Maximum Power Value       : 0.0 Watts
  Conservation Power Level  : 1=0.0W
- Avaya/Extreme Call Server(s): 69.26.36.53
- Avaya/Extreme IP Phone Address: 10.201.41.147 255.255.255.0
  Default Gateway Address   : 10.201.41.1
- Avaya/Extreme CNA Server: 0.0.0.0
- Avaya/Extreme File Server(s): 10.201.41.36
- Avaya/Extreme IEEE 802.1q Framing: Tagged

```

```
-----
LLDP Port 2:10 detected 1 neighbor
```

```

  Neighbor: 00:01:30:F9:9E:80/8:10, age 20 seconds
  - Chassis ID type: MAC address (4)
    Chassis ID       : 00:01:30:F9:9E:80
  - Port ID type: ifName (5)
    Port ID          : "8:10"
  - Time To Live: 120 seconds
  - System Description: "ExtremeXOS version 12.0.0.6 v1200b6 by release-ma\
                        nager on Mon Mar 19 00:43:19 PDT 2007"

```

History

This command was first available in ExtremeXOS 11.2. Information on the LLDP MED extension and Avaya-Extreme proprietary TLVs was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show lldp statistics

```
show lldp {port [all | <port_list>]} statistics
```

Description

Displays statistical counters related to the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The following counters are presented with the standard command (taken from the IEEE 802.1ab MIB definition):

- Last table change time: Last time an entry in the LLDP database was added, changed or deleted.
- Number of table inserts: The number of times the complete set of information advertised by a particular neighbor has been inserted into tables.
- Number of table deletes: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables.
- Number of table drops: The number of times the complete set of information advertised by a particular neighbor could not be stored in memory because of insufficient resources.
- Number of table age outs: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables because the information timeliness interval has expired.
- Tx Total: The number of LLDP frames transmitted by this switch on the indicated port.
- Tx Total Length Exceeded: The number of LLDP frames sent out on this port that could not hold all the information configured because the total frame length would exceed the maximum LDDPDU size of 1500 bytes.
- Rx Total: The number of valid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- Rx Discarded: The number of LLDP frames received by this switch on the indicated port, and then discarded for any reason.
- Rx Errors: The number of invalid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- TLVs Discarded: The number of LLDP TLVs discarded for any reason by this switch on the indicated port.
- TLVs Unrecognized: The number of LLDP TLVs received on the given port that are not recognized by the switch.

Example

The following example lists statistical counters for all ports on the switch:

```
# show lldp port all statistics
```

```
Last table change time   : Fri Dec 17 10:42:33 2004
Number of Table Inserts  : 3
Number of Table Deletes  : 0
Number of Table Drops    : 0
Number of Table Age Outs : 0
```

Port	Tx Total	Tx Length Exceeded	Rx Total	Rx Discarded	Rx Errors	TLVs Discarded	TLVs Unrecogn.
1:1	189	0	5654	0	0	0	0
2:2	188	0	565	0	0	0	0

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure lldp

```
unconfigure lldp {ports [all | <port_list>]}
```

Description

Leaves LLDP enabled and configured; restores the LLDP timer default values.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

When you issue the global `unconfigure lldp`, only the LLDP timers are reset to default values. All the configured TLVs remain on the ports remain, and LLDP remains enabled.

When you use the keyword `ports`, the TLVs for each port are returned to the five default TLVs. LLDP remains enabled.

Example

The following command restores LLDP factory default TLVs for ports 1:4 to 1:8:

```
unconfigure lldp ports 1:4 - 1:8
```

History

This command was first available in ExtremeXOS 11.2.

The keyword `port` was changed to `ports` in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

10 Power Over Ethernet Commands

Power over Ethernet (PoE) is an effective method of supplying 48 VDC power to certain types of powered devices (PDs) through Category 5 or Category 3 twisted pair Ethernet cables. PDs include wireless access points, IP telephones, laptop computers, web cameras, and other devices. With PoE, a single Ethernet cable supplies power and the data connection, reducing costs associated with separate power cabling and supply. PoE for ExtremeXOS includes a method of detection to assure that power is delivered to devices that meet the IEEE 802.3af specification for PoE, as well as to many legacy devices.

Extreme Networks PoE Devices

The following lists the Extreme Networks devices that support PoE and the minimum required software:

- G48P module for the BlackDiamond 8800 series switch—ExtremeXOS 11.1 and higher
- G48Pe module for the BlackDiamond 8800 series switch—ExtremeXOS 11.5 and higher
- Summit X450e-24p switch—ExtremeXOS 11.5 and higher
- Summit X450e-48p switch—ExtremeXOS 11.6 and higher

Summary of PoE Software Features

The Extreme Networks PoE devices support the following PoE software features:

- Configuration and control of the power distribution for PoE at the system, slot, and port levels
- Real-time discovery and classification of 802.3af-compliant PDs and many legacy (non-standard) devices
- Monitor and control of PoE fault conditions
- Support for configuring and monitoring PoE status at the system, slot, and port levels
- LED control for indicating the port's PoE inline power state
- Management of an over-subscribed power budget
- Beginning with ExtremeXOS version 11.3, support for hitless failover in a chassis with two MSMs
- Beginning with ExtremeXOS 12.0, support for failover in a SummitStack

For more conceptual information about configuring and managing PoE, see the *ExtremeXOS Concepts Guide*.

clear inline-power stats ports

```
clear inline-power stats ports [all | <port_list>]
```

Description

Clears the inline statistics for the selected port to zero.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Use this command to clear all the information displayed by the `show inline-power stats ports <port_list>` command.

Example

The following command clears the inline statistics for ports 1-8 on slot 3 on a modular switch:

```
clear inline-power stats ports 3:1-3:8
```

The following command displays cleared inline power configuration information for ports 1-8 in slot 3:

```
show inline-power stats ports 3:1-3:8
```

Following is sample output from this command:

Port	State	Class	STATISTICS				COUNTERS	
			Absent	InvSig	Denied	OverCurrent	Short	
3:1	delivering	class3	0	0	0	0	0	0
3:2	delivering	class3	0	0	0	0	0	0
3:3	searching	class0	0	0	0	0	0	0
3:4	searching	class0	0	0	0	0	0	0
3:5	searching	class0	0	0	0	0	0	0
3:6	searching	class0	0	0	0	0	0	0
3:7	searching	class0	0	0	0	0	0	0
3:8	searching	class0	0	0	0	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X450e-24p and X450e-48p switch (and on a SummitStack when these switches are present), and the Summit X250e-24p and X250e-48p switches.

configure inline-power budget

```
configure inline-power budget <num_watts> {slot <slot>}
```

Description

Sets the reserved power on the switch or specified slot to the specified watts.

Syntax Description

num_watts	Specifies the number of watts to reserve for specified switch or slot for inline power. Enter an integer. The minimum value is 37, or 0 if the slot is disabled; the maximum is 768; and the default value is 50.
slot	Specifies a slot. The slot must be configured to hold a PoE module.

Default

50 W.

Usage Guidelines

This command sets the budgeted power reserved for all PDs connected to the switch or specified slot in Watts. On a modular switch, none of the power budget on a specified slot can be used to power other slots or PDs on other slots. (refer to the *Extreme Networks Consolidated Hardware Guide* for complete information on power availability with this optional unit.)

On a modular switch, if you specify a slot that is not configured to hold a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

You can modify the power budget without disabling the switch or slot.

If the power consumption of the PDs on the switch or a specified slot exceeds this configured power budget, the system disconnects the lowest priority ports. (Refer to [configure inline-power priority ports](#) for information on configuring this parameter.)

If you attempt to configure this power budget for a value that the system cannot safely provide, the system returns an error message. To display inline power settings, use the command [show inline-power](#); to display the power for the entire switch, use the command [show power budget](#).



NOTE

You must disable inline power for the switch or the specified slot using the [disable inline-power slot](#) command prior to setting the budget to 0.

To reduce the chances of ports fluctuating between powered and non-powered states, newly inserted PDs are not powered when the actual delivered power for the module is within approximately 19 W of the configured inline power budget for that switch or slot. However, actual aggregate power can be

delivered up to the configured inline power budget for the switch or slot (for example, when delivered power from ports increases or when the configured inline power budget for the switch or slot is reduced).

For X450e-48p in a SummitStack, you may configure the PoE power budget per slot. The budget is limited to the power provided by the internal and external PSUs connected to the specific switch in the slot.

Example

The following command sets the power for slot 4 to 150 W on a modular switch:

```
configure inline-power budget 150 slot 4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, on a Summit X450e-48p, the Summit X250e-24p and X250e-48p switches, or on a SummitStack when this switch is present.

configure inline-power disconnect-precedence

```
configure inline-power disconnect-precedence [deny-port | lowest-priority]
```

Description

Configures the disconnect precedence priority for the switch when a new PD is detected and the measured inline power for that switch or specified slot is within 19 W of the switch's or slot's PoE power budget.

Syntax Description

deny-port	Specifies power be denied to PD requesting power, regardless of priority.
lowest-priority	Specifies power be withdrawn from lowest-priority port(s) when next PD requesting power connects.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the switch and for the entire modular switch; you cannot configure this per slot or per port.

If the power supplied to the PDs on a switch or specified slot exceeds the power that was budgeted for that switch or specified slot, the system disconnects power to one or more ports to prevent power overload. Refer to [configure inline-power budget](#) for information on configuring and modifying the power budgeted for each switch or specified slot.

You configure the switch to either deny power to the next PD that requests power on that switch or slot, regardless of the priority, or to disconnect those PDs on ports with lower priorities until there is enough power for the new PD. If you select this last argument and you did not configure port priorities or if several ports have the same priority, the switch withdraws power (or disconnects) those ports with the *highest* port number (s). Refer to [configure inline-power priority ports](#) for information on configuring the PoE priority for the ports.

The default value is deny-port. So, if you do not change the default value and the switch's or slot's power is exceeded, the next PD requesting power will not be connected.

When the setting is lowest priority, the switch continues dropping ports with the lowest configured PoE port priorities, or the highest port number in the case of equal PoE port priorities, until there is enough power for the requesting PD.

Example

The following command sets the switch to withdraw power from the lowest-priority port(s):

```
configure inline-power disconnect-precedence lowest-priority
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X450e-48p switch, the Summit X250e-24p and X250e-48p switches, and SummitStack when this switch is present.

configure inline-power label ports

```
configure inline-power label <string> ports <port_list>
```

Description

Lets you create your own label for a specified PoE port or group of PoE ports.

Syntax Description

string	Specifies a name up to 15 characters in length to identify the specified power port(s).
port_list	Specifies one or more ports or slots and ports.

Default

No label.

Usage Guidelines

Use the `show inline-power configuration ports` command, as shown in the following example, to display inline power configuration information, including the label (if any) for each port:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command on a modular switch:

Port	Config	Operator	Limit	Priority	Label
3:1	Enabled	16000	mW	Low	finance
3:2	Enabled	15000	mW	Low	finance
3:3	Enabled	15000	mW	Low	
3:4	Enabled	15000	mW	Low	
3:5	Enabled	15000	mW	Low	
3:6	Enabled	15000	mW	Low	marketing
3:7	Enabled	15000	mW	Low	marketing
3:8	Enabled	15000	mW	Low	marketing
3:9	Enabled	15000	mW	Low	
3:10	Enabled	15000	mW	Low	

Example

The following command assigns the name “alpha-test_1” to port 1 on slot 4:

```
config inline-power label alpha-test_1 ports 4:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, and the Summit X450e-24p and X450e-48p switches when these switches are present.

configure inline-power operator-limit ports

```
configure inline-power operator-limit <milliwatts> ports [all |<port_list>]
```

Description

Sets the power limit allowed for PDs connected to the specified ports.

Syntax Description

milliwatts	An integer specifying the maximum allowed power in milliwatts; the range is 3000 to 16800 mW. NOTE: If you attempt to enter a higher value, the switch returns an error message.
port_list	Specifies one or more ports or slots and ports.

Default

15400 mW.

Usage Guidelines

This command sets the power limit that a PD can draw on the specified ports. Range is 3000 to 16800 mW; the default value is 15400 mW.

If the measured power for a specified port exceeds the port's operator limit, the power is withdrawn from that port and the port moves into a fault state.

If you try to set an operator-limit outside the accepted range, the system returns the following error message:

```
Error: Invalid operator-limit value. Must be in the range of 3000-16800 mW
```

Example

The following command sets the limit for legacy PDs on ports 3 – 6 of slot 5 on a modular switch to 10000 mW:

```
configure inline-power operator-limit 10000 ports 5:3-5:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

configure inline-power priority ports

```
configure inline-power priority [critical | high | low] ports <port_list>
```

Description

Sets the PoE priority on the specified ports.

Syntax Description

critical high low	Sets the PoE priority for the specified ports.
port_list	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

The system allocates power to those ports with the highest priorities first. This command can also be used in conjunction with the [configure inline-power disconnect-precedence](#) command. If you configure the disconnect precedence as lowest priority, then newly detected PDs will be powered if that port has higher priority than the existing powered ports.

If there are multiple ports at the same priority level (either configured or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command assigns a critical PoE priority on ports 4 – 6 on slot 3 on a modular switch:

```
configure inline-power priority critical ports 3:4-3:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, and the Summit X450e-48p switch.

configure inline-power usage-threshold

```
configure inline-power usage-threshold <threshold>
```

Description

Sets the inline power usage SNMP event threshold.

Syntax Description

threshold	Specifies the percentage of budgeted power used on any PoE module or stand-alone switch that causes the system to send an SNMP event and create a log message. The range 1 to 99; the default value is 70.

Default

70.

Usage Guidelines

This command sets the threshold for generating an SNMP event and an Event Management System (EMS) message. On a modular switch, this threshold is when the measured power for a PoE module compared to the budgeted power for that slot exceeds a certain value. On stand-alone switches, this threshold applies to the total power available to the entire switch. The configured threshold value initiates the event and message once that percentage of the budgeted power is being used.

On a modular switch, the PoE threshold applies only to the percentage *per slot* of measured to budgeted power use; it does not apply systemwide.

The system generates an additional SNMP event and EMS message once the power usage falls *below* the threshold again; once the condition clears.

Example

The following command sets the inline power usage alarm threshold at 75%:

```
configure inline-power usage-threshold 75
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

disable inline-power

```
disable inline-power
```

Description

Shuts down PoE power currently provided on all ports on all slots.

Syntax Description

This command has no arguments or variables

Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. Using the `disable inline-power` command shuts down inline power currently provided on the entire switch or to specified ports and slots. Disabling inline power to a switch, port, or slot immediately removes power to any connected PDs. By default, inline power provided to all ports is enabled.



NOTE

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

On modular switches, disabling inline power does not allow PoE power reserved for slots to be allocated to other slots that may be needing more power to become operational. However, when you issue the command `disable slot` on a slot holding a PoE module, the inline power is also disabled; that slot is totally offline.



NOTE

Inline power cannot be delivered to connected PDs unless the Summit X450e-24p switch or Summit X450e-48p switch is powered; on modular switches, the BlackDiamond 8800 chassis and module must be powered.

Example

The following command shuts down inline power currently provided to all ports and all slots:

```
disable inline-power
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p and X450e-48p switches, and on a SummitStack when these switches are present.

disable inline-power legacy

```
disable inline-power legacy
```

Description

Disables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

Example

The following command disables capacitance detection of PDs on the switch:

```
disable inline-power legacy
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p and Summit X450e-48p switches, and on a SummitStack when these switches are present.

disable inline-power legacy slot

```
disable inline-power legacy slot <slot>
```

Description

Disables the non-standard (or capacitance) power detection mechanism for the specified slot.

Syntax Description

slot	Disables non-standard power detection for specified slot on a modular switch.

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch or specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command disables capacitance detection of PDs on slot 3 of a modular switch:

```
disable inline-power legacy slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

disable inline-power ports

```
disable inline-power ports [all | <port_list>]
```

Description

Shuts down PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Disables inline power to all ports on the switch.
port_list	Disables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to ports immediately removes power to any connected PDs. By default, the capability to provide inline power to all ports is enabled.



NOTE

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Disabling inline power to a port providing power to a PD immediately removes power to the PD.



NOTE

On a modular switch, PoE power removed from ports using this command can be used by other ports on the same module.

Example

The following command shuts down inline power currently provided to ports 4 and 5 on slot 3 on a modular switch:

```
disable inline-power ports 3:4-5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p and Summit X450e-48p switches, and on a SummitStack when these switches are present.

disable inline-power slot

```
disable inline-power slot <slot>
```

Description

Shuts down PoE power currently provided to the specified slot.

Syntax Description

slot	Disables inline power to specified slot.

Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, the capability to provide inline power to a slot is enabled.

Disabling a slot using this command does not change the power budgeted to a specified slot using the [configure inline-power budget](#) command; nor can that power be used by PDs connected to any other slot.



NOTE

You can set the reserved power budget to 0 for a slot if, and only if, you first issue this command.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command removes power to all PDs on slot 3:

```
disable inline-power slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, and on a SummitStack when a Summit X450e-48p switch is present.

enable inline-power

```
enable inline-power
```

Description

Enables PoE power to all ports; on a modular switch, this is all ports on all slots.

Syntax Description

This command has no arguments or variables.

Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. By default, inline power provided to all ports is enabled.

Enabling inline power starts the PoE detection process used to discover, classify, and power remote PDs.



NOTE

When you are working on a modular switch, if your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

However, when you issue the command `disable slot` for a modular switch on a slot holding a PoE module, the inline power is also disabled; that slot is totally offline.



NOTE

Inline power cannot be delivered to connected PDs unless the Summit X450e-24p switch or Summit X450e-48p switch is powered; in a modular switch, both the BlackDiamond 8800 chassis and module must be powered.

On modular switches and the Summit X450e-48p switch, to deliver inline power to connected PDs, you must also reserve adequate power for the switch or specified slot with the PDs using the `configure inline-power budget` command.

Example

The following command enables inline power currently provided to all ports and all slots:

```
enable inline-power
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

enable inline-power legacy

```
enable inline-power legacy
```

Description

Enables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and *after* an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



CAUTION

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

Example

The following command enables capacitance detection of PDs on the switch:

```
enable inline-power legacy
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit X250e-24p and X250e-48p switches, and the Summit X450e-24p switch and the Summit X450e-48p switch.

enable inline-power legacy slot

```
enable inline-power legacy slot <slot>
```

Description

Enables non-standard (or capacitance) power detection mechanism for the specified slot on a modular switch.

Syntax Description

slot	Enables non-standard power detection for specified slot on a modular switch.

Default

Disable.

Usage Guidelines

This command enables the non-standard power-detection mechanism on the specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and *after* an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



CAUTION

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

On stack, if you do not specify a slot number, the command operates on all active nodes. The command operates *only* on nodes in the active topology.

Example

The following command enables capacitance detection of PDs on slot 3 on a modular switch:

```
enable inline-power legacy slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

enable inline-power ports

```
enable inline-power ports [all | <port_list>]
```

Description

Enables PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Enables inline power to all ports on the switch.
port_list	Enables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to a port immediately removes power to any connected PD. By default, inline power provided to all ports is enabled.

On modular switches and the Summit X450e-48p switch, to deliver inline power to ports with connected PDs, you must also reserve power for the slot with the PDs using the [configure inline-power budget](#) command. If you do not have enough reserved power for the port, that port moves into a Denied state.



NOTE

On modular switches, if your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the [disable inline-power](#) command does *not* affect the data traffic traversing the port. And, disabling the port using the [disable port](#) command does not affect the inline power supplied to the port.

Example

The following command enables inline power to ports 4 and 5 on slot 3 on a modular switch:

```
enable inline-power ports 3:4-5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

enable inline-power slot

```
enable inline-power slot <slot>
```

Description

Enables PoE power to the specified slot on modular switches.

Syntax Description

slot	Enables inline power to specified slot.

Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, inline power provided to all slots is enabled.

To deliver inline power to slots, you must reserve power for that slot using the [configure inline-power budget](#) command. By default, each PoE module has 50 W of power reserved for inline power.



NOTE

If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the [disable inline-power](#) command does *not* affect the data traffic traversing the slot. And, disabling the slot using the [disable slot](#) command does not affect the inline power supplied to the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command makes inline power available to slot 3:

```
enable inline-power slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

reset inline-power ports

```
reset inline-power ports <port_list>
```

Description

Power cycles the specified ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports for which power is to be reset.

Default

N/A.

Usage Guidelines

This command power cycles the specified ports. Ports are immediately disabled and then re-enabled, allowing remote PDs to be power-cycled.

This command affects only inline power; it does not affect network connectivity for the port(s).

Example

The following command resets power for port 4 on slot 3 on a modular switch:

```
reset inline-power ports 3:4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

show inline-power

```
show inline-power
```

Description

Displays inline power status information for the specified PoE switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output varies depending on the PoE device you are using.

On the Summit X450e-24p and the Summit X450e-48p switch, the output indicates the following inline power status information:

- Inline power status—The status of inline power. The status conditions are:
 - Enabled
 - Disabled
- Power usage threshold
- Disconnect precedence
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational
 - Not operational
 - Disabled
 - Subsystem failure
- Measured power—The amount of power, in watts, that currently being used by the switch.
- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

On the Summit X450e-48p switch, the output indicates the following inline power status information:



NOTE

For additional information on inline power parameters, refer to the [show power budget](#) command.

Example

The following command displays inline power status for the switch:

```
show inline-power
```

Following is sample output from this command for the Summit X450e-24p switch:

```
Inline Power System Information
Configured           : Enabled
Power Usage Threshold : 70 percent
Disconnect Precedence : deny-port
                    Budgeted      Measured
Firmware Status      Power (Watts) Power (Watts)  Legacy
Operational          405 W         0 W           Disabled
```

Following is sample output from this command for the Summit X450e-24p switch and the Summit X450e-48p switch:

```
show inline-power
Inline Power System Information
Configured           : Enabled
Power Usage Threshold : 70 percent
Disconnect Precedence : deny-port
                    Budgeted      Measured
Firmware Status      Power (Watts) Power (Watts)  Legacy
Operational          405 W         397 W           Disabled
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch and the Summit X450e-48p switch.

show inline-power configuration ports

```
show inline-power configuration ports <port_list>
```

Description

Displays inline power configuration information for the specified ports.

Syntax Description

port_list	Specifies one or more ports.

Default

N/A.

Usage Guidelines

The output displays the following inline power configuration information for the specified ports:

- Config—Indicates whether the port is enabled to provide inline power:
 - Enabled: The port can provide inline power.
 - Disabled: The port cannot provide inline power.
- Operator Limit—Displays the configured limit, in milliwatts, for inline power on the port.
- Label—Displays a text string, if any, associated with the port.

The following also displays for this command on modular PoE devices and the Summit X450e-48p switch:

- Priority—Displays inline power priority of the port, which is used when the disconnect precedence is set to lowest priority:
 - Low
 - High
 - Critical

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on a modular switch:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command:

Port	Config	Operator	Limit	Priority	Label
3:1	Enabled	15000	mW	Low	
3:2	Enabled	15000	mW	Low	
3:3	Enabled	15000	mW	Low	
3:4	Enabled	15000	mW	Low	
3:5	Enabled	15000	mW	Low	
3:6	Enabled	15000	mW	Low	
3:7	Enabled	15000	mW	Low	
3:8	Enabled	15000	mW	Low	
3:9	Enabled	15000	mW	Low	
3:10	Enabled	15000	mW	Low	

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

show inline-power info ports

```
show inline-power info {detail} ports <port_list>
```

Description

Displays inline power information for the specified ports.

Syntax Description

port_list	Specifies one or more ports.

Default

N/A.

Usage Guidelines



NOTE

Ports in the denied or faulted state periodically display the searching state as the hardware retests the PD state.

You can use this command to generate a summary report or a detailed report.

Summary output displays the following inline power information for the specified ports:

- State—Displays the port power state:
 - Disabled
 - Searching
 - Delivering
 - Faulted
 - Disconnected
 - Other
 - Denied
- PD's power class—Displays the class type of the connected PD:
 - "----": disabled or searching
 - "class0": class 0 device
 - "class1": class 1 device
 - "class2": class 2 device
 - "class3": class 3 device
 - "class4": class 4 device
- Volts—Displays the measured voltage. A value from 0 to 2 is valid for ports that are in a searching state.

- Curr—Displays the measured current, in milliamperes, drawn by the PD.
- Power—Displays the measured power, in watts, supplied to the PD.
- Fault—Displays the fault value:
 - None
 - UV/OV fault
 - UV/OV spike
 - Over current
 - Overload
 - Undefined
 - Underload
 - HW fault
 - Discovery resistance fail
 - Operator limit violation
 - Disconnect
 - Discovery resistance, A2D failure
 - Classify, A2D failure
 - Sample, A2D failure
 - Device fault, A2D failure
 - Force on error

The detail command lists all inline power information for the selected ports. Detail output displays the following information:

- Configured Admin State—Displays the port's configured state; Enabled or Disabled.
- Inline Power State—Displays the port power state.
- MIB Detect Status—Displays the port state as reported by SNMP; valid values are as follows:
 - disabled
 - searching
 - delivering
 - fault
 - test
 - otherFault
 - denyLowPriority
- Label—Displays the port's configured label.
- Operator Limit—Displays the port's configured operator limit value.
- PD Class—Displays the class type of connected PD:
- Max Allowed Power—Displays the amount of maximum allowed power for a device of this class.
- Measured Power—Displays the measured power, in watts, supplied to the PD.
- Line Voltage—Displays the measured voltage. A value from 0 to 2 is valid for ports in a searching state.
- Current—Displays the measured current, in milliamperes, drawn by the PD.
- Fault Status—Displays the fault value.
- Detailed Status

The following information displays only with modular PoE devices and the Summit X450e-48p switch:

- Priority—Displays the port's configured PoE priority value, as follows:
 - Critical
 - High
 - Low

Example

The following command displays summary inline power information for ports 1 to 3 on slot 3 on a modular switch:

```
show inline-power info ports 3:1-3
```

Following is sample output from this command:

Port	State	Class	Volts	Curr (mA)	Power (Watts)	Fault
3:1	delivering	class3	48.3	192	9.300	None
3:2	delivering	class3	48.3	192	9.300	None
3:3	searching	-----	0.0	0	0.0	None

The following command displays detail inline power information for port 1 on slot 3:

```
show inline-power info detail port 3:1
```

Following is sample output from this command:

Port 3:1

```
Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status      : delivering
Label                  :
Operator Limit         : 16800 milliwatts
PD Class               : class3
Max Allowed Power      : 15.400 W
Measured Power         : 9.400 W
Line Voltage           : 48.3 Volts
Current                : 193 mA
Fault Status           : None
Detailed Status        :
Priority               : low
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

show inline-power slot

```
show inline-power slot <slot>
```

Description

Displays inline power information for the specified slot on modular switches.

Syntax Description

slot	Specifies the slot.

Default

N/A.

Usage Guidelines

On modular switches, the output indicates the following inline power status for each system:

- Configured power
 - Enabled
 - Disabled
- System power surplus
- Redundant power surplus
- Power usage threshold
- Disconnect precedence
- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

On modular switches, the output indicates the following inline power status information for each slot:

- Inline power status—The status of inline power. The status conditions are:
 - Enabled
 - Disabled
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational
 - Not operational
 - Disabled
 - Subsystem failure
 - Card not present
 - Slot disabled
- Budgeted power—The amount of power, in watts, that is available to the slot.
- Measured power—The amount of power, in watts, that currently being used by the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command displays inline power information for slot 3 on a modular switch:

```
show inline-power slot 3
```

Following is sample output from this command on a modular switch:

```

                Inline Power System Information
Configured      : Enabled
System Power Surplus      : 1500 Watts available for budgeting
Redundant Power Surplus   : 465 Watts available for budgeting to maintain N+1
Power Usage Threshold    : 70 percent (per slot)
Disconnect Precedence     : lowest-priority
Legacy Mode          : Disabled

Slot  Inline-Power  Firmware Status  Budgeted      Measured
      Inline-Power  Firmware Status  Power (Watts) Power (Watts)
3      Enabled      Operational      50 W          9 W
4      Enabled      Card Not Present  ( 50 W)       n/a
7      Enabled      Operational      50 W          0 W
Note: A budget value in parentheses is not allocated from the system power
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

show inline-power stats

```
show inline-power stats
```

Description

Displays inline power statistics for the specified switch.

Syntax Description

There are no variables or parameters for this command.

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the switch. Unlike the values displayed with the `show inline-power stats ports` command, these values are current readings; not cumulative counters.

Example

The following command displays inline power statistics information for the Summit X450e-24p switch:

```
show inline-power stats
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Firmware status           : Operational
Firmware revision         : 292b1

Total ports powered       : 7
Total ports awaiting power : 17
Total ports faulted       : 0
Total ports disabled      : 0
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch and the Summit X450e-48p switch.

show inline-power stats ports

```
show inline-power stats ports <port_list>
```

Description

Displays inline power statistics for the specified ports.

Syntax Description

port_list	Specifies one or more slots and ports.

Default

N/A.

Usage Guidelines

The output displays the following inline power statistics for the specified ports:

- State—Displays the port power state:
 - Disabled
 - Searching
 - Delivering
 - Faulted
 - Disconnected
 - Other
 - Denied
- PD's power class—Displays the class type of the connected PD:
 - "-----": disabled or searching
 - "class0": class 0 device
 - "class1": class 1 device
 - "class2": class 2 device
 - "class3": class 3 device
 - "class4": class 4 device
- Absent—Displays the number of times the port was disconnected.
- InvSig—Displays the number of times the port had an invalid signature.
- Denied—Displays the number of times the port was denied.
- Over-current—Displays the number of times the port entered an overcurrent state.
- Short—Displays the number of times the port entered undercurrent state.

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on a modular switch:

```
show inline-power stats ports 3:1-10
```

Following is sample output from this command:

Port	State	Class	STATISTICS				COUNTERS	
			Absent	InvSig	Denied	OverCurrent	Short	
3:1	delivering	class3	0	0	0	18	0	
3:2	delivering	class3	0	0	0	0	0	
3:3	searching	class0	0	0	0	0	0	
3:4	searching	class0	0	0	0	0	0	
3:5	searching	class0	0	0	0	0	0	
3:6	searching	class0	0	0	0	0	0	
3:7	searching	class0	0	0	0	0	0	
3:8	searching	class0	0	0	0	0	0	
3:9	searching	class0	0	0	0	0	0	
3:10	searching	class0	0	0	0	0	0	

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, and the Summit X450e-24p switch, and the Summit X450e-48p switch.

show inline-power stats slot

```
show inline-power stats slot <slot>
```

Description

Displays inline power statistics for the specified slot on modular switches.

Syntax Description

slot	Specifies the slot.

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the selected slots. Unlike the values displayed with the `show inline-power stats ports` command, these values (displayed with the `show inline-power stats slot` command) are current readings; not cumulative counters.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command displays inline power statistics information for slot 3 on a modular switch:

```
show inline-power stats slot 3
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Slot: 3
Firmware status      : Operational
Firmware revision    : 292b1

Total ports powered   : 7
Total ports awaiting power : 41
Total ports faulted   : 0
Total ports disabled  : 0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

unconfigure inline-power budget slot

```
unconfigure inline-power budget slot <slot>
```

Description

Unconfigures the inline reserved power on modular switches on the specified slot and returns the power budget on that slot to the default value of 50 W.

Syntax Description

slot	Specifies the slot.

Default

50 W.

Usage Guidelines

This command unconfigures any previously configured power budget for the specified slot and resets the budgeted power reserved for all PDs connected to this slot to 50 W. The rest of the previously configured power budget on this slot cannot be used to power other slots or PDs on other slots (unless you explicitly reconfigure the power budget for other slots).

If you specify a slot that does not have a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

Example

The following command resets the power for slot 4 to 50 W:

```
unconfigure inline-power budget slot 4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch.

unconfigure inline-power disconnect-precedence

```
unconfigure inline-power disconnect-precedence
```

Description

On a modular switch and the Summit X450e-48p switch, unconfigures the disconnect precedence setting and returns the switch to the default disconnect precedence value of deny port.

Syntax Description

This command has no arguments or variables.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the entire switch; you cannot configure this per slot or per port.

Unconfigures the PoE disconnect precedence previously set for the Summit X450e-48p switch or modular switch and returns the disconnect precedence to the default value of deny port. Deny port denies power to the next PD that requests inline power from the slot when the inline power budget for the switch or slot is reached, regardless of the inline power port priority.

Example

The following command resets the switch to the PoE disconnect precedence value, which is deny port:

```
unconfigure inline-power disconnect-precedence
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch the Summit X250e-24p and X250e-48p switches, and the Summit X450e-48p switch.

unconfigure inline-power operator-limit ports

```
unconfigure inline-power operator-limit ports [all |<port_list>]
```

Description

Unconfigures the PoE operator limit setting and resets the power limit allowed for PDs connected to the specified ports to the default value of 15400 mW.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more slots and ports.

Default

15400 mW.

Usage Guidelines

This command unconfigures any previously configured operator limit for the specified ports. It resets the maximum power that any PD can draw to 15400 mW.

Example

The following command resets the limit on ports 3 to 6 of slot 5 on a modular switch to the default value of 15400 mW:

```
unconfigure inline-power operator-limit ports 5:3-5:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

unconfigure inline-power priority ports

```
unconfigure inline-power priority ports [all | <port_list>]
```

Description

On modular switches and the Summit X450e-48p switch, unconfigures the PoE priority on the specified ports, and returns the ports to the default PoE port priority value of low.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

Use this to reset the PoE port priority on specified ports on modular switches and the Summit X450e-48p switch to the default value of low.

If there are multiple ports on the modular switch or the Summit X450e-48p switch at the same priority level (either configured, or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command resets the PoE priority on ports 4 – 6 on slot 3 to low:

```
unconfigure inline-power priority ports 3:4-3:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch the Summit X250e-24p and X250e-48p switches, and the Summit X450e-48p switch.

unconfigure inline-power usage-threshold

```
unconfigure inline-power usage-threshold
```

Description

Unconfigures the inline power usage alarm threshold and returns threshold to the default value of 70%.

Syntax Description

This command has no arguments or variables.

Default

70.

Usage Guidelines

This command unconfigures the PoE usage threshold setting for initiating SNMP event and EMS messages and returns the switch's inline power usage threshold for to 70%. The system initiates an event and message once that percentage of the budgeted power is being used.

On stand-alone switches, this PoE threshold applies to the *entire* switch. On modular switches, the threshold applies only to the percentage *per slot* of measured to budgeted power use; the threshold does *not* apply to the entire switch.

The system generates an additional SNMP event and EMS message once the power usage falls *below* the threshold again; once the condition clears.

Example

The following command resets the inline power usage alarm threshold to 70%:

```
unconfigure inline-power usage-threshold
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the G48P and G48Pe modules installed in the BlackDiamond 8800 series switch, the Summit X250e-24p and X250e-48p switches, the Summit X450e-24p switch, and the Summit X450e-48p switch.

This chapter describes commands for:

- Configuring and managing the Event Management System/Logging
- Configuring and monitoring system health and statistics
- Enabling, disabling, and configuring the Extreme Link Status Monitoring (ELSM) protocol
- Enabling and disabling the collection of remote monitoring (RMON) statistics on the switch
- Enabling, disabling, and configuring sFlow® statistics collection

Event Management System

When an event occurs on a switch, the Event Management System (EMS) allows you to send messages generated by these events to a specified log target. You can send messages to the memory buffer, NVRAM, the console display, the current session, to a syslog host, or to the other Management Switch Fabric Module (MSM). The log messages contain configuration and fault information pertaining to the device. You can format the log messages to contain various items of information, but typically a message consists of:

- Timestamp—The timestamp records when the event occurred.
- Severity level:
 - Critical—A desired switch function is inoperable. The switch may need to be reset.
 - Error—A problem is interfering with normal operation.
 - Warning—An abnormal condition exists that may lead to a function failure.
 - Notice—A normal but significant condition has been detected; the system is functioning as expected.
 - Info—Actions and events that are consistent with expected behavior.
 - Debug—Summary, Debug-Verbose, and Debug-Data—Information that is useful when performing detailed trouble shooting procedures.

By default, log entries that are assigned a critical, error, or warning level are considered static entries and remain in the NVRAM log target after a switch reboot.

- Component—The component refers to the specific functional area to which the error refers.
- Message—The message contains the log information with text that is specific to the problem.

The switch maintains a configurable number of messages in its internal (memory-buffer) log (1000 by default). You can display a snapshot of the log at any time. In addition to viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console display or telnet session. In addition to maintaining an internal log, the switch supports remote logging by way of the UNIX syslog host facility.

Beginning with ExtremeXOS 11.2, EMS supports IPv6 as a parameter for filtering events.

Extreme Link Status Monitoring

ExtremeXOS 11.4 introduces support for the Extreme Link Status Monitoring (ELSM) protocol. ELSM is an Extreme Networks proprietary protocol that monitors network health by detecting CPU and remote link failures. ELSM is available only on Extreme Networks devices and operates on a point-to-point basis. You configure ELSM on the ports that connect to other network devices and on both sides of the peer connection.

ELSM monitors network health by exchanging various hello messages between two ELSM peers. ELSM uses an open-ended protocol, which means that an ELSM-enabled port expects to send and receive hello messages from its peer. The Layer 2 connection between ports determines the peer connection. Peers can be either directly connected or separated by one or more hubs. If there is a direct connection between peers, they are considered neighbors.

If ELSM detects a failure, the ELSM-enabled port responds by blocking traffic on that port. For example, if a peer stops receiving messages from its peer, ELSM brings down that connection. ELSM does this by blocking all incoming and outgoing data traffic on the port and notifying applications that the link is down.

In some situations, a software or hardware fault may prevent the CPU from transmitting or receiving packets, thereby leading to the sudden failure of the CPU. If the CPU is unable to process or send packets, ELSM isolates the connections to the faulty switch from the rest of the network. If the switch fabric sends packets during a CPU failure, the switch may appear healthy when it is not. For example, if hardware forwarding is active and software forwarding experiences a failure, traffic forwarding may continue. Such failures can trigger control protocols such as Extreme Standby Router Protocol (ESRP) or Ethernet Automatic Protection Switching (EAPS) to select different devices to resume forwarding. This recovery action, combined with the CPU failure, can lead to loops in a Layer 2 network.

Configuring ELSM on Extreme Networks devices running ExtremeXOS is backward compatible with Extreme Networks devices running ExtremeWare.

sFlow Statistics

sFlow® is a technology for monitoring traffic in data networks containing switches and routers. It relies on statistical sampling of packets from high-speed networks, plus periodic gathering of the statistics. A User Datagram Protocol (UDP) datagram format is defined to send the information to an external entity for analysis. sFlow consists of a (Management Information Base) MIB and a specification of the packet format for forwarding information to a remote agent. Details of sFlow specifications can be found in RFC 3176 and at the following website:

<http://www.sflow.org>

ExtremeXOS allows you to collect sFlow statistics on a per port basis. An agent, residing locally on the switch, sends data to a collector that resides on another machine. You configure the local agent, the address of the remote collector, and the ports of interest for sFlow statistics gathering. You can also modify default values for how frequently on average a sample is taken, how often the data is sent to the collector, and the maximum load allowed on the CPU before throttling the statistics gathering.

Licensing

You need an Advanced Edge, Core, or Advanced Core license to configure and use the sFlow features described in this chapter.

sFlow functionality is not available with an Edge license.

The BlackDiamond 10808 switch with an MSM-1 module or an MSM1-XL module ships with a Core or Advanced Core license, respectively.

The BlackDiamond 12804 switch with an MSM-5 module or an MSM-5R module ships with a Core license. The BlackDiamond 12802 switch ships with an Advanced Edge license.

The BlackDiamond 8800 series switch with an MSM-G8X module or an MSM-48 module, the Summit X450 series switch, and the Summit X450a series switch ship with an Advanced Edge license.

The Summit X450e series switches ship with an Edge license, which does not support sFlow. You can, however, upgrade to an Advanced Edge license that supports sFlow.



NOTE

If you have a Summit X450e series switch, you must upgrade to an Advanced Edge license to configure and use the sFlow features described in this chapter.

For more information about software licensing, including how to obtain and upgrade your license, see [Chapter 2, “ExtremeXOS Overview,”](#) in the *ExtremeXOS Concepts Guide*.

RMON

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1757 and RFC 2021, which allows you to monitor LANs remotely.

Using the RMON capabilities of the switch allows network administrators to improve system efficiency and reduce the load on the network.

The IETF defines nine groups of Ethernet RMON statistics. The switch supports the following four of these groups, as defined in RFC 1757:

- Statistics
- History
- Alarms
- Events

The switch also supports the following parameters for configuring the RMON probe and the trap destination table, as defined in RFC 2021:

- probeCapabilities
- probeSoftwareRev
- probeHardwareRev

- probeDateTime
- probeResetControl
- trapDestTable

clear counters

```
clear counters
```

Description

Clears all switch statistics and port counters, including port packet statistics, bridging statistics, IP statistics, and log event counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics and port counters before you clear them. Use the `show ports` command to view port statistics. Use the `show log counters` command to show event statistics.

The CLI also provides a number of options that you can specify with the `clear counters` command. If you specify an option, the switch only clears the statistics for that option. For example, if you want to clear, reset only the STP statistics and counters, use the `clear counters stp` command. Please refer to the specific chapter in the *ExtremeXOS Command Reference Guide* for more detailed information about those commands.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears all switch statistics and port counters:

```
clear counters
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear elsm ports auto-restart

```
clear elsm ports <portlist> auto-restart
```

Description

Clears one or more ELSM-enabled ports that are in the Down-Stuck state.

Syntax Description

portlist	Specifies the ELSM-enabled ports that are permanently in the Down-Stuck state.

Default

N/A.

Usage Guidelines

If you do not have automatic restart enabled, use this command to transition ELSM-enabled ports that are permanently in the Down-Stuck state to the Down state. You can also use the [enable elsm ports <portlist> auto-restart](#) command to transition a port from the Down-Stuck state to the Down state.

For information about the ELSM-enabled ports states, see the command [show elsm ports on page 754](#).

If automatic restart is enabled (this is the default behavior), automatic restart automatically transitions the ports from the Down-Stuck state to the Down state. For more information, see the command [enable elsm ports auto-restart on page 735](#).

Example

The following command transitions the ports from the Down-Stuck state to the Down state:

```
clear elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

clear elsm ports counters

```
clear elsm {ports <portlist>} counters
```

Description

Clears the statistics gathered by ELSM for the specified ports or for all ports.

Syntax Description

portlist	Specifies the ELSM-enabled ports for which ELSM statistics are being cleared.

Default

N/A.

Usage Guidelines

You should view the ELSM statistics and counters before you clear them. To view ELSM-specific counter information, use the `show elsm ports <all | portlist>` command. To view summary ELSM information, including the ports configured for ELSM, use the `show elsm` command.

Use this command to clear only the ELSM-related counters. To clear all of the counters on the switch, including those related to ELSM, use the `clear counters` command.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counter, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears the statistics gathered by ELSM for slot 2, ports 1-2:

```
clear elsm ports 2:1-2:2 counters
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

clear log

```
clear log {error-led | static | messages [memory-buffer | nvram]}
```

Description

Clears the log messages in memory and NVRAM, and clears the ERR LED on the MSM.

Syntax Description

error-led	Clears the ERR LED on the MSM.
static	Specifies that the messages in the NVRAM and memory-buffer targets are cleared, and the ERR LED on the MSM is cleared.
memory-buffer	Clears entries from the memory buffer.
nvram	Clears entries from NVRAM.

Default

N/A.

Usage Guidelines

The switch log tracks configuration and fault information pertaining to the device.

By default, log entries that are sent to the NVRAM remain in the log after a switch reboot. The `clear log` and `clear log messages memory-buffer` commands remove entries in the memory buffer target; the `clear log static` and `clear log messages nvram` commands remove messages from the NVRAM target. In addition, the `clear log static` command will also clear the memory buffer target.

On modular switches and SummitStack, there are three ways to clear the ERR LED: clear the log, reboot the switch, or use the `clear log error-led` command. To clear the ERR LED without rebooting the switch or clearing the log messages, use the `clear log error-led` command.

Execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following command clears all log messages, from the NVRAM:

```
clear log static
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear log counters

```
clear log counters [<event-condition> | [all | <event-component>] {severity
<severity> {only}}]
```

Description

Clears the incident counters for events.

Syntax Description

event-condition	Specifies the event condition counter to clear.
all	Specifies that all events counters are to be cleared.
event-component	Specifies that all the event counters associated with a particular component should be cleared.
severity	Specifies the minimum severity level of event counters to clear (if the keyword only is omitted).
only	Specifies that only event counters of the specified severity level are to be cleared.

Default

If severity is not specified, then the event counters of any severity are cleared in the specified component.

Usage Guidelines

This command sets the incident counters to zero for each event specified. To display event counters, use the following command:

```
show log counters
```

See the command [show log](#) on [page 763](#) for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events {detail}
```

To get a listing of the components present in the system, use the following command:

```
show log components
```

In a SummitStack, execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following command clears the event counters for event conditions of severity error or greater in the component *BGP*:

```
clear log counters "BGP" severity error
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear sys-recovery-level

```
clear sys-recovery-level
```

Description

If configured and the switch detects a hardware fault and enters the shutdown state, this command clears the shutdown state and renders the switch, I/O, or MSM module(s) operational.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If you configure the switch or one or more modules to shutdown upon detecting a hardware fault, and the switch or module enters the shutdown state, you must explicitly clear the shutdown state and reset the switch or the affected modules for the switch to become operational.

To clear the shutdown state, use the following command:

```
clear sys-recovery-level
```

The switch prompts you to confirm this action. The following is a sample confirmation message:

```
Are you sure you want to clear sys-recovery-level? (y/n)
```

Enter y to confirm this action and clear the shutdown state. Enter n or press [Enter] to cancel this action.

Modular Switches Only. On a modular switch, after using the `clear sys-recovery-level` command, you must reset each affected module.

If you configured only a few I/O modules to shutdown, reset each affected I/O module as follows:

- 1 Disable the slot using the `disable slot <slot>` command.
- 2 Re-enable the slot using the `enable slot <slot>` command.



NOTE

You must complete this procedure for each module that enters the shutdown state.

If you configured all I/O modules or one or more MSMs to shutdown, use the `reboot` command to reboot the switch and reset all affected modules.

After you clear the shutdown state and reset the affected module, each port is brought offline and then back online before the module and the entire system is operational.

Summit Family of Switches Only. On a Summit family switch, after you clear the shutdown state, use the `reboot` command to bring the switch and ports back online. After you use the `reboot` command, the switch is operational.

Differences in the Command Line Prompt. When an exclamation point (!) appears in front of the command line prompt, it indicates that one or more slots of a modular switch or the entire stand-alone switch is shut down as a result of your hardware recovery configuration and a switch error.

The following is sample truncated output of the command line prompt for a modular switch:

```
The I/O modules in the following slots are shut down: 1,3
Use the "clear sys-recovery-level" command to restore I/O modules
```

```
! BD-8810.1 #
```

The following is sample output for the Summit family of switches:

```
All switch ports have been shut down.
Use the "clear sys-recovery-level" command to restore all ports.
```

```
! SummitX450-24x.1 #
```

Example

The following command clears the shutdown state:

```
clear sys-recovery-level
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure elsm ports hellotime

```
configure elsm ports <portlist> hellotime <hello_time>
```

Description

Configures the ELSM hello timer by specifying the time between consecutive hello messages for the specified ports.

Syntax Description

portlist	Specifies the port or ports for which the ELSM hello timer should be configured.
hello_time	Specifies the time in seconds between consecutive hello messages. Use the same value for the hello interval on peer ports. The default value is 1 second, and the range is 1 to 128 seconds.

Default

The default is 1 second.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages.

ELSM uses two types of hello messages to communicate the health of the network to other ELSM ports:

- Hello+ — The ELSM-enabled port receives a hello message from its peer and no problem is detected.
- Hello- — The ELSM-enabled port does not receive a hello message from its peer.

ELSM also has hello transmit states. The hello transmit states display the current state of transmitted ELSM hello messages. For more information about the hello transmit states, see the `show elsm ports` command.

A high hello timer value can increase the time it takes for the ELSM-enabled port to enter the Up state. The down timer is $(2 + \text{hold threshold}) * \text{hello timer}$. Assuming the default value of 2 for the hold threshold, configuring a hello timer of 128 seconds creates a down timer of $(2 + 2) 128$, or 512 seconds. In this scenario it would take 512 seconds for the port to transition from the Down to the Up state.

If you modify the hello timer on one port, Extreme Networks recommends that you use the same hello timer value on its peer port.

Example

The following command specifies 5 seconds between consecutive ELSM hello messages for slot 2, ports 1-2 on the switch:

```
configure elsm ports 2:1-2:2 hellotime 5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure elsm ports hold-threshold

```
configure elsm ports <portlist> hold-threshold <hold_threshold>
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Down-Wait state to the Up state.

Syntax Description

portlist	Specifies the port or ports for which the ELSM hold threshold should be configured.
hold_threshold	Specifies the number of Hello+ messages required to transition from the Down-Wait state to the Up state. The default is 2 messages, and the range is 1 to 40 messages.

Default

The default is 2 Hello+ messages.

Usage Guidelines

The port begins in the Down state, so the first received Hello+ message transitions the ELSM-enabled port from the Down state to the Down-Wait state. After that transition, the configured hold-threshold value determines the number of Hello+ messages required to transition from Down-Wait state to the Up state.

The ELSM hold threshold determines the number of Hello+ messages the ELSM peer port must receive to transition from the Down-Wait state to the Up state. For example, a threshold of 1 means the ELSM port must receive at least one Hello+ message to transition from the Down-Wait state to the Up state.

After the down timer expires, the port checks the number of Hello+ messages against the hold threshold. If the number of Hello+ messages received is greater than or equal to the configured hold threshold, the ELSM receive port moves from the Down-Wait state to the Up state.

If the number of Hello+ messages received is less than the configured hold threshold, the ELSM receive port moves from the Down-Wait state back to the Down state and begins the process again.

If you modify the hold threshold on one port, Extreme Networks recommends that you use the same hold threshold value on its peer port.

You configure the hold threshold on a per-port basis, not on a per-switch basis.

Example

The following command specifies that two Hello+ messages are required for the ELSM receive ports configured on slot 2, ports 1-2, to transition from the Down-Wait state to the Up state:

```
configure elsm hold-threshold 2 ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure elsm ports uptimer-threshold

```
configure elsm ports <portlist> uptimer-threshold <uptimer_threshold>
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Up state to the Down state.

Syntax Description

portlist	Specifies the port or ports for which the ELSM hold threshold should be configured.
uptimer_threshold	Specifies the number of Hello+ messages required to transition from the Up-state to the Down state. The default is 6messages, and the range is 3 to 60 messages.

Default

The default is 6 Hello+ messages.

Usage Guidelines

The ELSM up timer begins when the ELSM-enabled port enters the UP state. Each time the port receives a Hello+ message, the timer restarts. Up timer is Uptimer_threshold * hello timer. When the Up timer expires, it transits from UP state to DOWN state.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure log display

```
configure log display <severity> {only}
```

Description

Configures the real-time log-level message to display.

Syntax Description

severity	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.
only	Specifies only log messages of the specified severity level.

Default

If not specified, messages of all severities are displayed on the console display.

Usage Guidelines

You must enable the log display before messages are displayed on the log display. Use the [enable log display](#) command to enable the log display. This allows you to configure the system to maintain a running real-time display of log messages on the console.

Severity filters the log to display messages with the selected severity or higher (more critical). Severities include critical, error, warning, info, notice, debug-summary, debug-verbose, and debug-data.

You can also control log data to different targets. The command equivalent to [configure log display](#) is the following:

```
configure log target console-display severity <severity>
```

To display the current configuration of the log display, use the following command:

```
show log configuration target console-display
```

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Example

The following command configures the system log to maintain a running real-time display of log messages of critical severity or higher:

```
configure log display critical
```

The following command configures the system log to maintain a running real-time display of only log messages of critical severity:

```
configure log display critical only
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure log filter events

```
configure log filter <name> [add | delete] {exclude} events [<event-  
condition> | [all | <event-component>] {severity <severity> {only}}]
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

name	Specifies the filter to configure.
add	Add the specified events to the filter
delete	Remove the specified events from the filter
exclude	Events matching the specified events will be excluded
event-condition	Specifies an individual event.
all	Specifies all components and subcomponents.
event-component	Specifies all the events associated with a particular component.
severity	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.

Default

If the `exclude` keyword is not used, the events will be included by the filter. If `severity` is not specified, then the filter will use the component default severity threshold (see the note [on page 665](#) when `delete` or `exclude` is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events. If you want to configure a filter to include or exclude incidents based on event parameter values (for example, MAC address or BGP Neighbor) see the command `configure log filter events match` [on page 667](#).

When the `add` keyword is used, the specified event name is added to the beginning of the filter item list maintained for this filter. The new filter item either includes the events specified, or if the `exclude` keyword is present, excludes the events specified.

The `delete` keyword is used to remove events from the filter item list that were previously added using the `add` command. All filter items currently in the filter item list that are identical to, or a subset of, the set of events specified in the `delete` command will be removed.

Event Filtering Process. From a logical standpoint, the filter associated with each enabled log target is examined to determine whether a message should be logged to that particular target. The determination is made for a given filter by comparing the incident with the most recently configured filter item first. If

the incident matches this filter item, the incident is either included or excluded, depending on whether the `exclude` keyword was used. Subsequent filter items on the list are compared if necessary. If the list of filter items has been exhausted with no match, the incident is excluded.

Events, Components, and Subcomponents. As mentioned, a single event can be included or excluded by specifying the event's name. Multiple events can be added or removed by specifying an ExtremeXOS component name plus an optional severity. Some components, such as *BGP*, contain subcomponents, such as *Keepalive*, which is specified as *BGP.Keepalive*. Either components or subcomponents can be specified. The keyword `all` in place of a component name can be used to indicate all ExtremeXOS components.

Severity Levels. When an individual event name is specified following the `events` keyword, no severity value is needed since each event has pre-assigned severity. When a component, subcomponent, or the `all` keyword is specified following the `events` keyword, a severity value is optional. If no severity is specified, the severity used for each applicable subcomponent is obtained from the pre-assigned severity threshold levels for those subcomponents. For example, if *STP* were specified as the component, and no severity is specified for the add of an include item, then only messages with severity of `error` and greater would be passed, since the threshold severity for the *STP* component is `error`. If *STP.InBPDU* were specified as the component, and no severity is specified, then only messages with severity of `warning` and greater would be passed, since the threshold severity for the *STP.InBPDU* subcomponent is `warning`. Use the `show log components` command to see this information.

The severity keyword `all` can be used as a convenience when `delete` or `exclude` is specified. The use of `delete` (or `exclude`) with severity `all` deletes (or excludes) previously added events of the same component of all severity values.



NOTE

If no severity is specified when `delete` or `exclude` is specified, severity `all` is used

If the `only` keyword is present following the severity value, then only the events in the specified component at that exact severity are included. Without the `only` keyword, events in the specified component at that severity or more urgent are included. For example, using the option `severity warning` implies `critical`, `error`, or `warning` events, whereas the option `severity warning only` implies `warning` events only. Severity `all only` is not a valid choice.

Any EMS events with severity `debug-summary`, `debug-verbose`, or `debug-data` will not be logged unless debug mode is enabled

Filter Optimization. Each time a `configure log filter` command is issued for a given filter name, the events specified are compared against the current configuration of the filter to try to logically simplify the configuration.

For example, if the command:

```
configure log filter bgpFilter1 add events bgp.keepalive severity error only
```

were to be followed by the command:

```
configure log filter bgpFilter1 add events bgp severity info
```

the filter item in the first command is automatically deleted since all events in the *BGP.Keepalive* subcomponent at severity `error` would be also included as part of the second command, making the first command redundant.

More Information. See the command `show log` on page 763 for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command adds all STP component events at severity `info` to the filter *mySTPFilter*:

```
configure log filter myStpFilter add events stp severity info
```

The following command adds the *STP.OutBPDU* subcomponent, at the pre-defined severity level for that component, to the filter *myStpFilter*:

```
configure log filter myStpFilter add events stp.outbpdu
```

The following command excludes one particular event, *STP.InBPDU.Drop*, from the filter:

```
configure log filter myStpFilter add exclude events stp.inbpdu.drop
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure log filter events match

```
configure log filter <name> [add | delete] {exclude} events [<event-  
condition> | [all | <event-component>] {severity <severity> {only}}] [match  
| strict-match] <type> <value>
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events and match parameter values.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

name	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the filter will be excluded.
event-condition	Specifies the event condition.
all	Specifies all events.
event-component	Specifies all the events associated with a particular component.
severity	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.
match	Specifies events whose parameter values match the <type> <value> pair.
strict-match	Specifies events whose parameter values match the <type> <value> pair, and possess all the parameters specified.
type	Specifies the type of parameter to match. For more information about types and values see “Types and Values” on page 668 .
value	Specifies the value of the parameter to match. For more information about types and values see “Types and Values” on page 668 .

Default

If the `exclude` keyword is not used, the events will be included by the filter. If `severity` is not specified, then the filter will use the component default severity threshold (see the note on [on page 665](#) when `delete` or `exclude` is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events that match a list of <type> <value> pairs. This command is an extension of the command [configure log filter events](#), and adds the ability to filter incidents based on matching specified event parameter values to the event.

See the `configure log filter events` command on page 664 for more information on specifying and using filters, on event conditions and components, and on the details of the filtering process. The discussion here is about the concepts of matching `<type> <value>` pairs to more narrowly define filters.

Types and Values. Each event in ExtremeXOS is defined with a message format and zero or more parameter types. The `show log events` command on page 780 can be used to display event definitions (the event text and parameter types). The syntax for the parameter types (represented by `<type>` in the command syntax above) is:

```
[address-family [ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast]
| bgp-neighbor <ip address>
| bgp-routerid <ip address>
| eaps <eaps domain name>
| {destination | source} [ipaddress <ip address> | L4-port | mac-address ]
| esrp <esrp domain name>
| {egress | ingress} [slot <slot number> | ports <portlist>]
| ipaddress <ip address>
| L4-port <L4-port>
| mac-address <mac_address>
| netmask <netmask>
| number <number>
| port <portlist>
| process <process name>
| slot <slotid>
| string <match expression>
| vlan <vlan name>
| vlan tag <vlan tag>]
```



NOTE

The slot parameters are available only on modular switches.

Beginning with ExtremeXOS 11.2, you can specify the `ipaddress` type as IPv4 or IPv6, depending on the IP version. The following examples show how to configure IPv4 addresses and IPv6 addresses:

- IPv4 address

To configure an IP address, with a mask of 32 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.1
```

To configure a range of IP addresses with a mask of 8, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.0/8
```

- IPv6 address

To configure an IPv6 address, with a mask of 128 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::1
```

To configure a range of IPv6 addresses with a mask of 16, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::/16
```

- IPv6 scoped address

IPv6 scoped addresses consist of an IPv6 address and a VLAN. The following examples identify a link local IPv6 address.

To configure a scoped IPv6 address, with a mask of 128 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::1%Default
```


To configure a range of scoped IPv6 addresses with a mask of 16, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::/16%Default
```

To configure a scoped IPv6 address with any VLAN, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::/16%*
```

To configure any scoped IPv6 address with a specific VLAN, use the following command:

```
configure log filter myFilter add events all match ipaddress ::/0%Default
```



NOTE

In the previous example, if you specify the VLAN name, it must be a full match; wild cards are not allowed.

The <value> depends on the parameter type specified. As an example, an event may contain a physical port number, a source MAC address, and a destination MAC address. To allow only those incidents with a specific source MAC address, use the following in the command:

```
configure log filter myFilter add events aaa.radius.requestInit secerity notice match
source mac-address 00:01:30:23:C1:00
configure log filter myFilter add events bridge severity notice match source mac-
address 00:01:30:23:C1:00
```

The string type is used to match a specific string value of an event parameter, such as a user name. A string can be specified as a simple regular expression.

Match Versus Strict-Match. The `match` and `strict-match` keywords control the filter behavior for incidents whose event definition does not contain all the parameters specified in a `configure log filter events match` command. This is best explained with an example. Suppose an event in the XYZ component, named `XYZ.event5`, contains a physical port number, a source MAC address, but no destination MAC address. If you configure a filter to match a source MAC address and a destination MAC address, `XYZ.event5` will match the filter when the source MAC address matches regardless of the destination MAC address, since the event contains no destination MAC address. If you specify the `strict-match` keyword, then the filter will never match, since `XYZ.event5` does not contain the destination MAC address.

In other words, if the `match` keyword is specified, an incident will pass a filter so long as all parameter values in the incident match those in the match criteria, but all parameter types in the match criteria need not be present in the event definition.

More Information. See the command `show log` on page 763 for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

By default, all log targets are associated with the built-in filter, *DefaultFilter*. Therefore, the most straightforward way to send additional messages to a log target is to modify *DefaultFilter*. In the following example, the command modifies the built-in filter to allow incidents in the *STP* component, and all subcomponents of *STP*, of severity critical, error, warning, notice and info. For any of these events containing a physical port number as a match parameter, limit the incidents to only those occurring on physical ports 3, 4 and 5 on slot 1, and all ports on slot 2:

```
configure log filter DefaultFilter add events stp severity info match ports 1:3-1:5,  
2:*
```

If desired, issue the `unconfigure log DefaultFilter` command to restore the *DefaultFilter* back to its original configuration.

History

This command was first available in ExtremeXOS 10.1.

New parameter <type> values, including `esrp` and `eaps` were added in ExtremeXOS 11.0 and 11.1.

Support for IPv6 addresses was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure log target filter

```
configure log target [console | memory-buffer | primary-msm | primary-node
| backup-msm | backup-node | nvram | session | syslog [all | <ipaddress> |
<ipPort> {vr <vr_name>} [local0 ... local7]]] filter <filter-name>
{severity <severity> {only}}
```

Description

Associates a filter to a target.

In a stack, this command is applicable only to Master and Backup nodes. This command is not applicable to standby nodes.

Syntax Description

target	Specifies the device to send the log entries.
console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
primary-msm	Specifies the primary MSM.
primary-node	NOTE: This parameter is available only on modular switches. Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
backup-node	NOTE: This parameter is available only on modular switches. Specifies the backup node in a stack.
nvram	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog remote server.
all	Specifies all of the syslog remote servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
filter-name	Specifies the filter to associate with the target.
severity	Specifies the minimum severity level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

If severity is not specified, the severity level for the target is left unchanged. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command associates the specified filter and severity with the specified target. A filter limits messages sent to a target.

Although each target can be configured with its own filter, by default, all targets are associated with the built-in filter, *DefaultFilter*. Each target can also be configured with its own severity level. This provides the ability to associate multiple targets with the same filter, while having a configurable severity level for each target.

A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified. By default, the memory buffer, NVRAM, primary MSM, and backup MSM targets are enabled. For other targets, use the command [enable log target](#) on page 740.

[Table 12](#) describes the default characteristics of each type of target.

Table 12: Default target log characteristics

Target	Enabled	Severity Level
console display	no	info
memory buffer	yes	debug-data
NVRAM	yes	warning
primary MSM	yes	warning
backup MSM	yes	warning
session	no	info
syslog	no	debug-data

The built-in filter, *DefaultFilter*, and a severity level of *info* are used for each new telnet session. These values may be overridden on a per-session basis using the [configure log target filter](#) command and specify the target as *session*. Use the following form of the command for per-session configuration changes:

```
configure log target session filter <filter name> {severity <severity> {only}}
```

Configuration changes to the current session target are in effect only for the duration of the session, and are not saved in FLASH memory. The *session* option can also be used on the console display, if the changes are desired to be temporary. If changes to the console-display are to be permanent (saved to FLASH memory), use the following form of the command:

```
configure log target console filter <filter name> {severity <severity> {only}}
```

Modular Switches Only. If the condition for the *backup-msm* target is met by a message generated on the primary, the event is sent to the backup MSM. When the backup MSM receives the event, it will see if any of the local targets (*nvr*, *memory*, or *console*) are matched. If so it gets processed. The *session* and *syslog* targets are disabled on the backup MSM, as they are handled on the primary. If the

condition for the `primary-msm` target is met by a message generated on the backup, the event is sent to the primary MSM.

Note that the `backup-msm` target is only active on the primary MSM, and the `primary-msm` target is only active on the backup MSM.

SummitStack only. The `backup-node` target is only active on the primary-node, and the `primary-node` target is active on backup-node and standby-nodes.

Example

The following command sends log messages to the previously syslog host at 10.31.8.25, port 8993, and facility `local3`, that pass the filter `myFilter` and are of severity `warning` and above:

```
configure log target syslog 10.31.8.25:8993 local3 filter myFilter severity warning
```

The following command sends log messages to the current session, that pass the filter `myFilter` and are of severity `warning` and above:

```
configure log target session filter myFilter severity warning
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target format

```
configure log target [console | memory-buffer | nvram | session | syslog
[all | <ipaddress> | <ipPort>] {vr <vr_name>} {local0 ... local7}]]
format [timestamp [seconds | hundredths | none]
| date [dd-mm-yyyy | dd-Mmm-yyyy | mm-dd-yyyy | Mmm-dd | yyyy-mm-dd | none]
| severity
| event-name [component | condition | none | subcomponent]
| host-name
| priority
| process-name
| process-slot
| source-line
```

Description

Configures the formats of the displayed message, on a per-target basis.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
timestamp	Specifies a timestamp formatted to display seconds, hundredths, or none.
date	Specifies a date formatted as specified, or none.
severity	Specifies whether to include the severity.
event-name	Specifies how detailed the event description will be. Choose from none, component, subcomponent, or condition.
host-name	Specifies whether to include the syslog host name.
priority	Specifies whether to include the priority.
process-name	Specifies whether to include the internal process name.
process-slot	Specifies which slot number the message was generated. NOTE: This parameter is available only on modular switches.
source-line	Specifies whether to include the source file name and line number.

Default

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- severity—on
- event-name—condition
- host-name—off
- priority—off
- process-name—off
- process-slot—off (modular switches only)
- source-line—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd
- severity—on
- event-name—none
- host-name—off
- priority—on
- process-name—off
- process-slot—off (modular switches only)
- source-line—off

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command configures the format of the items that make up log messages. You can choose to include or exclude items and set the format for those items, but you cannot vary the order in which the items are assembled.

When applied to the targets `console` or `session`, the format specified is used for the messages sent to the console display or telnet session. Configuration changes to the `session` target, be it either a telnet or console display target session, are in effect only for the duration of the session, and are not saved in FLASH.

When this command is applied to the target `memory-buffer`, the format specified is used in subsequent `show log` and `upload log` commands. The format configured for the internal memory buffer can be overridden by specifying a format on the `show log` and `upload log` commands.

When this command is applied to the target `syslog`, the format specified is used for the messages sent to the specified syslog host.

Timestamps. Timestamps refer to the time an event occurred, and can be output in either seconds as described in RFC 3164 (for example, "13:42:56"), hundredths of a second (for example, "13:42:56.98"), or suppressed altogether. To display timestamps as hh:mm:ss, use the `seconds` keyword, to display as hh:mm:ss.HH, use the `hundredths` keyword, or to suppress timestamps altogether, use the `none` keyword. Timestamps are displayed in hundredths by default.

Date. The date an event occurred can be output as described in RFC 3164. Dates are output in different formats, depending on the keyword chosen. The following lists the `date` keyword options, and how the date "March 26, 2005" would be output:

- `Mmm-dd`—Mar 26
- `mm-dd-yyyy`—03/26/2005
- `dd-mm-yyyy`—26-03-2005
- `yyyy-mm-dd`—2005-03-26
- `dd-Mmm-yyyy`—26-Mar-2005

Dates are suppressed altogether by specifying `none`. Dates are displayed as `mm-dd-yyyy` by default.

Severity. A four-letter abbreviation of the severity of the event can be output by specifying `severity` on or suppressed by specifying `severity off`. The default setting is `severity on`. The abbreviations are: Crit, Erro, Warn, Noti, Info, Summ, Verb, and Data. These correspond to: Critical, Error, Warning, Notice, Informational, Debug-Summary, Debug-Verbose, and Debug-Data.

Event Names. Event names can be output as the component name only by specifying `event-name component` and as component and subcomponent name with condition mnemonic by specifying `event-name condition`, or suppressed by specifying `event-name none`. The default setting is `event-name condition` to specify the complete name of the events.

Host Name. The configured SNMP name of the switch can be output as `HOSTNAME` described in RFC 3164 by specifying `host-name`. The default setting is off.

Process Name. For providing detailed information to technical support, the (internal) ExtremeXOS task names of the applications detecting the events can be displayed by specifying `process-name`. The default setting is off.

Process Slot. For providing detailed information to technical support, the slot from which the logged message was generated can be displayed by specifying `process-slot`. The default setting is off. This is available only on modular switches.

Process ID. For providing detailed information to technical support, the (internal) ExtremeXOS task identifiers of the applications detecting the events can be displayed by specifying `process-id`. The default setting is off.

Source Line. For providing detailed information to technical support, the application source file names and line numbers detecting the events can be displayed by specifying `source-line`. The default setting is off. You must enable debug mode using the `enable log debug-mode` command to view the source line information. For messages generated prior to enabling debug mode, the source line information is not displayed.

Example

In the following example, the switch generates the identical event from the component SNTP, using three different formats.

Using the default format for the session target, an example log message might appear as:

```
05/29/2005 12:15:25.00 <Warn:SNTP.RslvSrvrFail> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

If you set the current session format using the following command:

```
configure log target session format timestamp seconds date mm-dd-yyyy event-name
component
```

The same example would appear as:

```
05/29/2005 12:16:36 <Warn:SNTP> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

To provide some detailed information to technical support, you set the current session format using the following command:

```
configure log target session format timestamp hundredths date mmm-dd event-name
condition source-line process-name
```

The same example would appear as:

```
May 29 12:17:20.11 SNTP: <Warn:SNTP.RslvSrvrFail> tSntp: (sntpLib.c:606) The SNTP
server parameter value (TheWrongServer.example.com) can not be resolved.
```

History

This command was first available in ExtremeXOS 10.1.

The `ipPort` and `host-name` parameters were first introduced in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target match

```
configure log target [console | memory-buffer | nvram | primary-msm |
primary-node| backup-msm | backp-node | session | syslog [all | <ipaddress>
| <ipPort> {vr <vr_name>} [local0 ... local7]]] match [any |<match-
expression>]
```

Description

Associates a match expression to a target.

In a stack, this command is applicable only on a Master and Backup nodes. This command is not applicable for standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
	NOTE: This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
	NOTE: This parameter is available only on modular switches.
backup-node	Specifies the backup-node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
any	Specifies that any messages will match. This effectively removes a previously configured match expression.
match-expression	Specifies a regular expression. Only messages that match the regular expression will be sent.

Default

By default, targets do not have a match expression. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command configures the specified target with a match expression. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command [show log](#) on page 763 for a detailed description of simple regular expressions. By default, targets do not have a match expression.

Specifying `any` instead of `match-expression` effectively removes a match expression that had been previously configured, causing any message to be sent that has satisfied all of the other requirements.

To see the configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {<ipaddress> | <ipPort> |
vr <vr_name>} {[local0 ... local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command sends log messages to the current session, that pass the current filter and severity level, and contain the string `user5`:

```
configure log target session match user5
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target severity

```
configure log target [console | memory-buffer | nvram | primary-msm |
primaryr-node | backup-msm | backup-node | session | syslog [all |
<ipaddress> | <ipPort> {vr <vr_name>} [local0 ... local7]]] {severity
<severity> {only}}
```

Description

Sets the severity level of messages sent to the target.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
	NOTE: This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
	NOTE: This parameter is available only on modular switches.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
severity	Specifies the least severe level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

By default, targets are sent messages of the following severity level and above:

- console display—info
- memory buffer—debug-data
- NVRAM—warning
- session—info

- syslog—debug-data
- primary MSM—warning (modular switches only)
- backup MSM—warning (modular switches only)
- primary node—warning (stack only)
- backup node—warning (stack only)

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command configures the specified target with a severity level. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command [show log on page 763](#) for a detailed description of severity levels.

To see the current configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {<ipaddress> | <ipPort> |
vr <vr_name>} {[local0 ... local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command sends log messages to the current session, that pass the current filter at a severity level of info or greater, and contain the string *user5*:

```
configure log target session severity info
```

History

This command was first available in ExtremeXOS 10.1.

The *primary-msm* and *backup-msm* options were first available in ExtremeXOS 11.0.

The *ipPort* parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target syslog

```
configure log target syslog [all | <ipaddress> | <ipPort>] {vr <vr_name>}
{local0 ... local7} from <source-ip-address>
```

Description

Configures the syslog server's IP address for one or all syslog targets.

Syntax Description

syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog server's IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
source-ip-address	Specifies the local source IP address to use.

Default

If a virtual router is not specified, the following virtual routers are used:

- ExtremeXOS 10.1—VR-0
- ExtremeXOS 11.0 and later—VR-Mgmt

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Use this command to identify and configure the syslog server's IP address. By configuring a source IP address, the syslog server can identify from which switch it received the log message.

Options for configuring the remote syslog server include:

- **all**—Specifies all of the remote syslog server hosts.
- **ipaddress**—The IP address of the remote syslog server host.
- **ipPort**—The UDP port.
- **vr_name**—The virtual router that can reach the syslog host.
- **local0-local7**—The syslog facility level for local use.
- **from**—The local source IP address.

If you do not configure a source IP address for the syslog target, the switch uses the IP address in the configured VR that has the closed route to the destination.

Example

The following command configures the IP address for the specified syslog target named *orange*:

```
configure log target syslog orange from 10.234.56.78
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ports monitor vlan

```
configure ports [<port_list>|all] monitor vlan <vlan_name>
```

Description

Starts counting VLAN statistics on a port or a group of ports.



NOTE

This command is available only on the BlackDiamond 12800 series switches.

Syntax Description

port_list	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
vlan_name	Specifies a VLAN name.

Default

N/A

Usage Guidelines

When per-port monitoring is configured, the `show ports` command displays the latest statistics directly from the hardware in real time. This information is not logged.

Examples

The following command configures per-port monitoring for a set of ports on slot 8 for the VLAN named *accounting*:

```
configure ports 8:1-8:6 monitor vlan accounting
```

You can monitor up to 1000 VLANs on the same port by issuing the command multiple times. For example, if you want to monitor VLANs *dog1*, *dog2*, *dog3*, and *dog4* on slot 1, use the following commands:

```
configure ports 1:* monitor vlan dog1
configure ports 1:* monitor vlan dog2
configure ports 1:* monitor vlan dog3
configure ports 1:* monitor vlan dog4
```

After you have configured the ports for monitoring, you can use the `show ports {<port_list>} vlan statistics {no-refresh}` command to display information for the configured ports:

```
show ports 1:* vlan statistics {no-refresh}
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 12800 series platforms.

configure sflow agent ipaddress

```
configure sflow agent {ipaddress} <ip-address>
```

Description

Configures the sFlow agent's IP address.

Syntax Description

ip-address	Specifies the IP address from which sFlow data is sent on the switch.

Default

The default configured IP address is 0.0.0.0, but the effective IP address is the management port IP address.

Usage Guidelines

This command allows you to configure the IP address of the sFlow agent. Typically, you would set this to the IP address used to identify the switch in the network management tools that you use. The agent address is stored in the payload of the sFlow data, and is used by the sFlow collector to identify each agent uniquely. The default configured value is 0.0.0.0, but the switch will use the management port IP address if it exists.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow agent` will reset the agent parameter to the default.

Example

The following command sets the sFlow agent's IP address to 10.2.0.1:

```
configure sflow agent ipaddress 10.2.0.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sflow collector ipaddress

```
configure sflow collector {ipaddress} <ip-address> {port <udp-port-number>}
{vr <vrname>}
```

Description

Configures the sFlow collector IP address.

Syntax Description

ip-address	Specifies the IP address to send the sFlow data.
udp-port-number	Specifies the UDP port to send the sFlow data.
vrname	Specifies from which virtual router to send the sFlow data.
NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.	

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—VR-Mgmt (previously called VR-0).

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command allows you to configure where to send the sFlow data. You must specify an IP address for the sFlow data collector, and you may specify a particular UDP port, if your collector uses a non-standard port. You may also need to specify from which virtual router to send the data.

You can configure up to four sFlow collectors. Each unique IP address/UDP port/virtual router combination identifies a collector.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow collector` will reset the collector parameters to the default.

Example

The following command specifies that sFlow data should be sent to port 6343 at IP address 192.168.57.1 using the virtual router *VR-Mgmt*:

```
configure sflow collector ipaddress 192.168.57.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sflow max-cpu-sample-limit

```
configure sflow max-cpu-sample-limit <rate>
```

Description

Configures the maximum number of sFlow samples handled by the CPU per second.

Syntax Description

rate	Specifies the maximum sFlow samples per second.

Default

The default value is 2000 samples per second.

Usage Guidelines

This command configures the maximum number of samples sent to the CPU per second. If this rate is exceeded, the internal sFlow CPU throttling mechanism kicks in to limit the load on the CPU.

Every time the limit is reached, the sample rate is halved (the value of number in the `configure sflow sample-rate <number>` or `configure sflow ports <portlist> sample-rate <number>` command is doubled) on the slot (modular switch) or ports (stand-alone switch) on which maximum number of packets were received during the last snapshot.

This effectively halves the sampling frequency of all the ports on that slot or stand-alone switch with a sub-sampling factor of 1. The sampling frequency of ports on that slot or stand-alone switch with a sub-sampling factor greater than 1 will not change; the sub-sampling factor is also halved so that the same rate of samples are sent from that port.

The maximum CPU sample rate is based on the total number of samples received from all the sources. The valid range is 100 to 5000 samples per second.

Example

The following command specifies that the sFlow maximum CPU sample rate should be set to 4000 samples per second:

```
configure sflow max-cpu-sample-limit 4000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sflow poll-interval

```
configure sflow poll-interval <seconds>
```

Description

Configures the sFlow counter polling interval.

Syntax Description

seconds	Specifies the number of seconds between polling each counter. The value can range from 0 to 3600 seconds.

Default

The default polling interval is 20 seconds.

Usage Guidelines

Each sFlow statistics counter is polled at regular intervals, and this data is then sent to the sFlow collector. This command is used to set the polling interval. To manage CPU load, polling for sFlow enabled ports are distributed over the polling interval, so that all ports are not polled at the same instant. For example, if the polling interval is 20 seconds and there are twenty counters, data is collected successively every second.

Specifying a poll interval of 0 (zero) seconds disables polling.

Example

The following command sets the polling interval to 60 seconds:

```
configure sflow poll-interval 60
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sflow ports sample-rate

```
configure sflow ports <portlist> sample-rate <number>
```

Description

Configures the sFlow per-port sampling rate.

Syntax Description

portlist	Specifies a list of ports.
number	Specifies the fraction (1/number) of packets to be sampled.

Default

The default number is 8192, unless modified by the `configure sflow sample-rate` command.

Usage Guidelines

This command configures the sampling rate on a particular set of ports, and overrides the system-wide value set in the `configure sflow sample-rate` command. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 1 to 536870912.

BlackDiamond 10808 Switch Only. At the hardware level, the sampling rate for all ports on a slot is the same. Ports that are configured for fewer samples than other ports on the slot (the `number` parameter is larger) still take the same amount of samples at the hardware level as ports configured for more samples. The excess samples are discarded. This mechanism is displayed in the `show sflow {configuration}` command as the sub-sampling factor. For example, if one port is configured for a sample rate of 1024, and a second port on the same slot is configured for a sample rate of 4096, the second port's hardware actually samples one out of 1024 packets, but only saves the information for one out of four of those samples. In this example, the sub-sampling factor is four.

Configuring a lower number for the sample rate means that more samples will be taken, increasing the load on the switch. Do not configure the sample rate to a number lower than the default unless you are sure that the traffic rate on the source is low.

If the rate of samples going to the CPU exceeds a limit, set with the `configure sflow max-cpu-sample-limit` command, then the sample rate number will be doubled, effectively reducing the rate by one-half. Under those conditions, the configured rate may not be the actual rate of sampling. The output of the `show sflow configuration` command will display the configured and actual rate of sampling.

Summit Family of Switches and BlackDiamond 8800 a-Series and e-Series Modules. All ports on the switch or same I/O module are sampled individually.

Example

The following command sets the sample rate for the ports 4:6 to 4:10 to one packet out of every 16384:

```
configure sflow ports 4:6-4:10 sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sflow sample-rate

```
configure sflow sample-rate <number>
```

Description

Configures the sFlow default sampling rate.

Syntax Description

number	Specifies the fraction (1/number) of packets to be sampled.

Default

The default number is 8192.

Usage Guidelines

This command configures the default sampling rate. This is the rate that newly enabled sFlow ports will have their sample rate set to. Changing this rate will not affect currently enabled sFlow ports. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 1 to 536870912.

Configuring a lower number for the sample rate means that more samples will be taken, increasing the load on the switch. Do not configure the sample rate to a number lower than the default unless you are sure that the traffic rate on the source is low.

Summit X450a and X450e Series Switches and BlackDiamond 8800 a-Series and e-Series Modules Only. The minimum rate that these platforms sample is 1 out of every 256 packets. If you configure a rate to be less than 256, the switch automatically rounds up the sample rate to 256.

Example

The following command sets the sample rate to one packet out of every 16384:

```
configure sflow sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure sys-health-check all level

```
configure sys-health-check all level [normal | strict]
```

Description

Configures how ExtremeXOS handles faults for the BlackDiamond 8800 series and the Summit family of switches.

Syntax Description

normal	Upon a fault detection, the switch only sends a message to the syslog. This is the default setting.
strict	Upon a fault detection, the switch takes the action configured by the <code>configure sys-recovery-level slot</code> or the <code>configure sys-recovery-level switch</code> command.

Default

The default setting is normal.

Usage Guidelines

On the BlackDiamond 8800 series switch, use this command in conjunction with the `configure sys-recovery-level slot <slot_number> [none | reset | shutdown]` command to implement your network's fault handling strategy.

On the Summit family of switches, use this command in conjunction with the `configure sys-recovery-level switch [none | reset | shutdown]` command to implement your network's fault handling strategy.

ExtremeXOS 11.5 enhances the number of switch-fabric tests completed and monitored by the polling module of the system health checker. Additionally with ExtremeXOS 11.5, you can now configure how ExtremeXOS handles a detected fault based on the configuration of the `configure sys-recovery-level slot <slot_number> [none | reset | shutdown]` or the `configure sys-recovery-level switch [none | reset | shutdown]` command.

If you configure the `strict` parameter, the switch takes the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, which can include logging only or restarting, rebooting, or shutting down the suspect device.

To maintain a smooth upgrade for devices running ExtremeXOS 11.4 and earlier, the switch-fabric tests introduced in ExtremeXOS 11.5 are set to only log error messages ('normal mode') by default. However, Extreme Networks recommends that you configure 'strict mode' so the system can attempt to recover by utilizing the action configured in the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command (which by default is reset).

System Behavior for the BlackDiamond 8800 Series Switch. Depending on your switch configuration, [Table 13](#) shows how the BlackDiamond 8800 series switch behaves when ExtremeXOS 11.5 detects a fault:

Table 13: System behavior for the BlackDiamond 8800 series switch

Fault Handling Configuration	Module Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level slot none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level slot reset</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level slot none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot reset</code>	ExtremeXOS reboots the affected switch or module.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	ExtremeXOS shuts down the affected switch or module.

System Behavior for the Summit Family of Switches. Depending on your switch configuration, [Table 13](#) shows how the Summit family of switches behave when ExtremeXOS 11.5 detects a fault:

Table 14: System behavior for the Summit family of switches

Fault Handling Configuration	Hardware Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level switch none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level switch reset</code>	Same as above.
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level switch none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level switch reset</code>	ExtremeXOS reboots the affected switch.
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	ExtremeXOS shuts down the affected switch.

Displaying the System Health Check Setting. To display the system health check setting, including polling and how ExtremeXOS handles faults on the switch, use the following command:

```
show switch
```

The system health check setting, displayed as `SysHealth check`, shows the polling setting and how ExtremeXOS handles faults. The polling setting appears as `Enabled`, and the fault handling setting appears in parenthesis next to the polling setting. In the following truncated output from a BlackDiamond 8810 switch, the system health check setting appears as `SysHealth check: Enabled (Normal)`:

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
```

System MAC: 00:04:96:1F:A2:60

SysHealth check: Enabled (Normal)

Recovery Mode: None

System Watchdog: Enabled

If you use the `strict` parameter, which configures the switch to take the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, (Strict) would appear next to Enabled.

Example

On the BlackDiamond 8800 series switch, the following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level slot` command:

```
configure sys-health-check all level strict
```

On the Summit family of switches, the following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level switch` command:

```
configure sys-health-check all level strict
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 8800 series, SummitStack, and the Summit family of switches.

configure sys-health-check interval

```
configure sys-health-check interval <interval>
```

Description

Configures the frequency of sending backplane diagnostic packets and the polling interval.

Syntax Description

interval	BlackDiamond 10808 and BlackDiamond 12800 series switches—Specifies the frequency of sending backplane diagnostic packets and polling the ASIC counters. • If backplane diagnostic packets are enabled on a particular slot, the default value for sending diagnostic packets and polling the ASIC counters is 6 seconds on that slot. • If only polling occurs (this is the system default), the default value is 60 seconds. The polling interval is <code>interval*10</code> seconds. BlackDiamond 8800 series switch—Specifies the frequency of sending backplane diagnostic packets. • If backplane diagnostic packets are enabled on a particular slot, the default value for sending diagnostic packets is 5 seconds on that slot. • If only polling occurs (this is the system default), the default value is 5 seconds. (The polling interval is not a user-configured parameter, and polling always occurs.)

Default

Depending upon your platform, the following defaults apply:

- BlackDiamond 10808 and BlackDiamond 12800 series switches:
 - If backplane diagnostics packets are enabled on a particular slot, the default for sending packets and for polling is 6 seconds on that slot.
 - If only polling occurs, the default polling interval is 60 seconds (`interval*10` seconds).
- BlackDiamond 8800 series switch:
 - If backplane diagnostics are enabled on a particular slot, the default for sending packets is 5 seconds on that slot.
 - The polling interval is always 5 seconds (this is a not a user-configured parameter).

Usage Guidelines

Use this command with the guidance of Extreme Networks Technical Support personnel.

The system health checker tests I/O modules and the backplane by forwarding backplane diagnostic packets. On the BlackDiamond 10808 and the BlackDiamond 12800 series switches, additional checking for the validity of these packets is completed by performing a checksum. Use this command to configure the amount of time it takes for the packets to be forwarded and returned to the MSM.

To enable backplane diagnostic packets, use the `enable sys-health-check slot <slot>` command. With backplane diagnostic packets enabled on a specific slot, the `interval` option of the `configure sys-health-check interval` command specifies the frequency of sending backplane diagnostic packets. For example, if you specify an interval of 9, backplane diagnostic packets are sent every 9 seconds on only the enabled slot.

**NOTE**

Extreme Networks does not recommend configuring an interval of less than the default interval. Doing this can cause excessive CPU utilization.

BlackDiamond 10808 and BlackDiamond 12800 Series Switches Only. If backplane diagnostic packets are disabled, you can use this command to configure the global polling interval on the switch.

When you enable backplane diagnostic packets with the `enable sys-health-check slot <slot>` command, the `interval` option also specifies the polling frequency on that slot. For example, if you specify an interval of 9, backplane diagnostic packets are sent and polling occurs every 9 seconds for only the enabled slot.

You can also use the `configure sys-health-check interval` command to modify the global polling interval on the switch. With only polling enabled, the `interval` option specifies the frequency of polling on all slots. Therefore, if you configure the interval to 8 seconds, the polling interval is updated to 80 seconds (`interval*10` seconds).

If backplane diagnostic packets are enabled on a slot and you modify the polling interval on the switch, the new interval value takes precedence. For example, if you configure the interval to 7 seconds, the polling interval is updated to 70 seconds on all slots with backplane diagnostics disabled. On the slot with backplane diagnostic packets enabled, backplane diagnostic packets are sent and polling occurs every 7 seconds.

To return to the default of sending diagnostic packets every 6 seconds and the system polling every 60 seconds, specify 6 for the `interval`.

By default, the system health checker always polls and tracks the ASIC counters that collect correctable and uncorrectable packet memory errors, checksum errors, and parity errors on a per ASIC basis. You must enable the backplane diagnostic packets feature to send backplane diagnostic packets.

BlackDiamond 8800 Series Switch Only. By default, the system health checker always polls the control plane health between MSMs and I/O modules, monitors memory levels on the I/O module, monitors the health of the I/O module, and checks the health of applications and processes running on the I/O module. If the system health checker detects an error, the health checker notifies the MSM.

You must enable the backplane diagnostic packets feature to send backplane diagnostic packets. If you enable this feature, the system health checker tests the data link for a specific I/O module every 5 seconds by default. The MSM sends and receives diagnostic packets from the I/O module to determine the state and connectivity. If you disable backplane diagnostics, the system health checker stops sending backplane diagnostic packets.

Example

The following examples assume that you enabled backplane diagnostic packets on a specific I/O slot.

On the BlackDiamond 10808 or the BlackDiamond 12800 series switches, the following command configures the backplane diagnostic packet interval to 8 seconds and the polling interval to 80 seconds:

```
configure sys-health-check interval 8
```

On the BlackDiamond 8800 series switch, the following command configures the backplane diagnostic packet interval to 8 seconds:

```
configure sys-health-check interval 8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

configure sys-recovery-level

```
configure sys-recovery-level [all | none]
```

Description

Configures a recovery option for instances where a software exception occurs in ExtremeXOS.

Syntax Description

all	Configures ExtremeXOS to log an error into the syslog and reboot the system after any software task exception occurs.
none	Configures the recovery level to none. No action is taken when a software task exception occurs; there is no system reboot, which can cause unexpected switch behavior. NOTE: Use this parameter only under the guidance of Extreme Networks Technical Support personnel.

Default

The default setting is `all`.

Usage Guidelines

If the software fails, the switch automatically reboots or leaves the system in its current state. You must specify one of the following parameters for the system to respond to software failures:

- `all`—The system will send error messages to the syslog and reboot if any software task exception occurs.
- `none`—No action is taken when a software task exception occurs. The system does not reboot, which can cause unexpected switch behavior.



NOTE

Use the `none` parameter only under the guidance of Extreme Networks Technical Support personnel.

If the system recovery level is `all`, the system always reboots after a software task exception. This is the default setting and behavior, and Extreme Networks strongly recommends using the default setting.

Displaying the System Recovery Setting. To display the software recovery setting on the switch, use the following command:

```
show switch
```


This command displays general switch information, including the software recovery level. The following truncated output from a Summit X450 series switch displays the software recovery setting (displayed as `Recovery Mode`):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:20:B4:13

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
```

**NOTE**

All platforms display the software recovery setting as `Recovery Mode`.

Example

The following command configures a switch to not take an action when any software task exception occurs:

```
configure sys-recovery-level none
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sys-recovery-level slot

```
configure sys-recovery-level slot <slot_number> [none | reset | shutdown]
```

Description

Configures a recovery option for instances where an exception occurs on the specified MSM or I/O module.

Syntax Description

slot_number	Specifies the slot of the MSM or I/O module. - A and B—Indicate an MSM 1 through 10—Indicate an I/O module
none	Configures the MSM or I/O module to maintain its current state regardless of the detected hardware fault. The offending MSM or I/O module is not reset. For more information about the states of an MSM or I/O module see the show slot command.
reset	Configures the offending MSM or I/O module to reset upon a hardware fault detection. For more detailed information, see the “Usage Guidelines” described below.
shutdown	Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs remain operational for debugging purposes only. ExtremeXOS logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Default

The default setting is `reset`.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure the MSMs or I/O modules installed in a modular switch to take no action, automatically reset, shutdown, or if dual MSMs are installed, failover to the other MSM, if the switch detects a faulty MSM or I/O module. This enhanced level of recovery detects faults in the ASICs as well as packet buses.

You must specify one of the following parameters for the system to respond to MSM or I/O module failures:

- none**—Configures the MSM or I/O module to maintain its current state regardless of the detected fault. The offending MSM or I/O module is not reset. ExtremeXOS logs fault and error messages to the syslog and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
- reset**—Configures the offending MSM or I/O module to reset upon fault detection. ExtremeXOS logs fault, error, system reset, and system reboot messages to the syslog.
- shutdown**—Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs remain operational for debugging purposes

only. ExtremeXOS logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Depending on your configuration, the switch resets the offending MSM or I/O module if fault detection occurs. An offending MSM is reset any number of times, and the MSM is not permanently taken offline. On the BlackDiamond 10808 and BlackDiamond 12804 switches, an offending I/O module is reset a maximum of three times. On the BlackDiamond 8800 series switch, an offending I/O module is reset a maximum of five times. After the maximum number of resets, the I/O module is permanently taken offline.

Messages Displayed. If you configure the hardware recovery setting to either none (ignore) or shutdown, the switch prompts you to confirm this action. The following is a sample shutdown message:

```
Are you sure you want to shutdown on errors? (y/n)
```

Enter *y* to confirm this action and configure the hardware recovery level. Enter *n* or press [Enter] to cancel this action.

Taking Ports Offline. Beginning with ExtremeXOS 11.5, you can configure the switch to shut down one or more modules upon fault detection by specifying the `shutdown` option. If you configure one or more slots to shut down and the switch detects a hardware fault, all ports in all of the configured shut down slots are taken offline in response to the reported errors. (MSMs are available for debugging purposes only.)

The affected module remains in the shutdown state across additional reboots or power cycles until you explicitly clear the shutdown state. If a module enters the shutdown state, the module actually reboots and the `show slot` command displays the state of the slot as `Initialized`; however, the ports are shut down and taken offline. For more information about clearing the shutdown state, see the `clear sys-recovery-level` command.

Module Recovery Actions—BlackDiamond 8800 Series Switch Only. Table 15 describes the actions module recovery takes based on your module recovery setting. For example, if you configure a module recovery setting of `reset` for an I/O module, the module is reset a maximum of five times before it is taken permanently offline.

From left to right, the columns display the following information:

- **Module Recovery Setting**—This is the parameter used by the `configure sys-recovery-level slot` command to distinguish the module recovery behavior.
- **Hardware**—This indicates the hardware that you may have in your switch.
- **Action Taken**—This describes the action the hardware takes based on the module recovery setting.

Table 15: Module recovery actions for the BlackDiamond 8800 series switch

Module Recovery Setting	Hardware	Action Taken
none	Single MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	Dual MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.

Table 15: Module recovery actions for the BlackDiamond 8800 series switch (Continued)

Module Recovery Setting	Hardware	Action Taken
reset	I/O Module	The I/O module remains powered on in its current state. The switch sends error messages to the log and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	Single MSM	Resets the MSM.
	Dual MSM	Resets the primary MSM and fails over to the backup MSM.
shutdown	I/O Module	Resets the I/O module a maximum of five times. After the fifth time, the I/O module is permanently taken offline.
	Single MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	Dual MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	I/O Module	Reboots the I/O module. When the module comes up, the ports remain inactive because you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the I/O module to become operational. After you clear the shutdown state, you must reset each affected I/O module or reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.

Module Recovery Actions—BlackDiamond 10808 and BlackDiamond 12804 Switches Only. Table 16 describes the actions module recovery takes based on your module recovery and software recovery settings. For example, if you configure a module recovery setting of `reset` and a system recovery setting of `all` for an I/O module, the module is reset a maximum of three times before it is taken permanently offline.

From left to right, the columns display the following information:

- **Module Recovery Setting**—This is the parameter used by the `configure sys-recovery-level slot` command to distinguish the module recovery behavior.
- **System Recovery Setting**—This is the parameter used by the `configure sys-recovery-level` command to distinguish the software recovery behavior.
- **Hardware**—This lists the hardware that you may have in your switch
- **Action Taken**—This describes the action the hardware takes based on the module recovery and software recovery settings.

Table 16: Module recovery actions for the BlackDiamond 10808 and the BlackDiamond 12804 switches

Module Recovery Setting	System Recovery Setting	Hardware	Action Taken
reset	all	Single MSM	Resets the MSM.
		Dual MSM	Resets the primary MSM and fails over to the backup MSM.
		I/O Module	Resets the I/O module a maximum of three times. After the third time, the I/O module is permanently taken offline.
reset	none	Single MSM	Resets the MSM.
		Dual MSM	Fails over to the backup MSM.
		I/O Module	Resets the I/O module.
shutdown	all	Single MSM	The MSM is available for debugging purposes only; however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
		Dual MSM	The MSMs are available for debugging purposes only; however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
		I/O Module	Reboots the I/O module. When the module comes up, the ports remain inactive because you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the I/O module to become operational. After you clear the shutdown state, you must reset each affected I/O module or reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
shutdown	none	Single MSM	The MSM is available for debugging purposes only.
		Dual MSM	The MSMs are available for debugging purposes only.

Table 16: Module recovery actions for the BlackDiamond 10808 and the BlackDiamond 12804 switches (Continued)

Module Recovery Setting	System Recovery Setting	Hardware	Action Taken
		I/O Module	Reboots the I/O module. When the module comes up, the ports remain inactive because you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the I/O module to become operational. After you clear the shutdown state, you must reset each affected I/O module or reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.

Displaying the Module Recovery Setting. To display the module recovery setting, use the following command:

```
show slot
```

Beginning with ExtremeXOS 11.5, the `show slot` output has been modified to include the shutdown configuration. If you configure the module recovery setting to shutdown, the output displays an “E” flag that indicates any errors detected on the slot disables all ports on the slot. The “E” flag appears only if you configure the module recovery setting to shutdown.

**NOTE**

If you configure one or more slots for shut down and the switch detects a hardware fault on one of those slots, all of the configured slots enter the shutdown state and remain in that state until explicitly cleared.

If you configure the module recovery setting to none, the output displays an “e” flag that indicates no corrective actions will occur for the specified MSM or I/O module. The “e” flag appears only if you configure the module recovery setting to none.

The following sample output displays the module recovery action. In this example, notice the flags identified for slot 10:

Slots	Type	Configured	State	Ports	Flags
Slot-1	G48P	G48P	Operational	48	MB S
Slot-2	G24X	G24X	Operational	24	MB S
Slot-3	G48T	G48T	Operational	48	MB S
Slot-4			Empty	0	
Slot-5	G8X	G8X	Operational	8	MB S
Slot-6	G8X	G8X	Operational	8	MB S
Slot-7			Empty	0	
Slot-8	G48Te	G48Te	Operational	48	MB S
Slot-9	G48Ta		Operational	48	MB S
Slot-10	G48Pe	G48Pe	Operational	48	MB S E
MSM-A	MSM-G8X		Operational	0	S
MSM-B	MSM-G8X		Operational	0	S

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")
 e - Errors on slot will be ignored (no corrective action initiated)
 E - Errors on slot will disable all ports on slot

**NOTE**

In ExtremeXOS 11.4 and earlier, if you configure the module recovery setting to none, the output displays an "E" flag that indicates no corrective actions will occur for the specified MSM or I/O module. The "E" flag appears only if you configure the module recovery setting to none.

Displaying Detailed Module Recovery Information. To display the module recovery setting for a specific port on a module, including the current recovery mode, use the following command:

```
show slot <slot>
```

In addition to the information displayed with show slot, this command displays the module recovery setting configured on the slot. The following truncated output displays the module recovery setting (displayed as Recovery Mode) for the specified slot:

```
Slot-10 information:
  State:                Operational
  Download %:           100
  Flags:                MB S E
  Serial number:        800158-00-01 06014-00022
  Hw Module Type:       G48Pe
  SW Version:           11.5.0.4
  SW Build:             v1150b4
  Configured Type:      G48Pe
  Ports available:      48
  Recovery Mode:        Shutdown
```

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")
 e - Errors on slot will be ignored (no corrective action initiated)
 E - Errors on slot will disable all ports on slot

Troubleshooting Module Failures. If you experience an I/O module failure, use the following troubleshooting methods when you can bring the switch offline to solve or learn more about the problem:

- Restarting the I/O module—Use the `disable slot <slot>` command followed by the `enable slot <slot>` command to restart the offending I/O module. By issuing these commands, the I/O module and its associated fail counter is reset. If the module does not restart, or you continue to experience I/O module failure, please contact Extreme Networks Technical Support.
- Running diagnostics—Use the `run diagnostics normal <slot>` command to run operational diagnostics on the offending I/O module to ensure that you are not experiencing a hardware issue. If the module continues to enter the failed state, please contact Extreme Networks Technical Support.

If you experience an MSM failure, please contact Extreme Networks Technical Support.

Example

The following command configures a switch to not take an action if a hardware fault occurs:

```
configure sys-recovery-level slot none
```

History

This command was first available in ExtremeXOS 11.3.

The `shutdown` parameter was added in ExtremeXOS 11.5.

Platform Availability

This command is available only on modular switches.

configure sys-recovery-level switch

```
configure sys-recovery-level switch [none | reset | shutdown]
```

Description

Configures a recovery option for instances where a hardware exception occurs on the Summit family of switches.

Syntax Description

none	Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown. ExtremeXOS logs fault and error messages to the syslog.
reset	Configures the switch to reboot upon detecting a hardware fault. ExtremeXOS logs fault, error, system reset, and system reboot messages to the syslog.
shutdown	Configures the switch to shut down upon detecting a hardware fault. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only. If the switch shuts down, it remains in this state across additional reboots or power cycles until you explicitly clear the shutdown state.

Default

The default setting is `reset`.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure the Summit family of switches to take no action, automatically reboot, or shutdown if the switch detects a hardware fault. This enhanced level of recovery detects faults in the CPU.

You must specify one of the following parameters for the switch to respond to hardware failures:

- `none`—Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown.
- `reset`—Configures the switch to reboot upon detecting a hardware fault.
- `shutdown`—Configures the switch to shutdown upon fault detection. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only.

Messages Displayed. If you configure the hardware recovery setting to either `none` (ignore) or `shutdown`, the switch prompts you to confirm this action by displaying a message similar to the following:

```
Are you sure you want to shutdown on errors? (y/n)
```

Enter `y` to confirm this action and configure the hardware recovery level. Enter `n` or press [Enter] to cancel this action.

Displaying the Hardware Recovery Setting. To display the hardware recovery setting, use the following command:

```
show switch
```

If you change the hardware recovery setting from the default (reset) to either none (ignore) or shutdown, the Recovery Mode output is expanded to include a description of the hardware recovery mode. If you keep the default behavior or return to reset, the Recovery Mode output lists only the software recovery setting.

The following truncated output from a Summit X450 series switch displays the software recovery and hardware recovery settings (displayed as `Recovery Mode`):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A5:71
Recovery Mode:    All, Ignore
System Watchdog:  Enabled
```

If you configure the hardware recovery setting to none, the output displays “Ignore” to indicate that no corrective actions will occur on the switch. “Ignore” appears only if you configure the hardware recovery setting to none.

If you configure the hardware recovery setting to shutdown, the output displays “Shutdown” to indicate that the switch will shutdown if fault detection occurs. “Shutdown” appears only if you configure the hardware recovery setting to shutdown.

If you configure the hardware recovery setting to reset, the output only displays the software recovery mode.

Example

The following command configures the switch to not take an action if a hardware fault occurs:

```
configure sys-recovery-level switch none
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the Summit family of switches.

configure syslog add

```
configure syslog {add} [<ipaddress> | <ipPort>] {vr <vr_name>} [local0 ...
local7] {<severity>}
```

Description

Configures the remote syslog server host address, and filters messages to be sent to the remote syslog target.

Syntax Description

ipaddress	Specifies the remote syslog server IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
severity	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.

Default

If a severity level is not specified, all messages are sent to the remote syslog server target. If a virtual router is not specified, VR-Mgmt is used. If UDP port is not specified, 514 is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Options for configuring the remote syslog server include:

- ipaddress—The IP address of the remote syslog server host.
- ipPort—The UDP port.
- local0-local7—The syslog facility level for local use.
- vr_name—The virtual router that can reach the syslog host.
- severity—Filters the messages sent to the remote syslog server target to have the selected severity or higher (more critical). Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.

The switch log overwrites existing log messages in a wrap-around memory buffer, which may cause you to lose valuable information once the buffer becomes full. The remote syslog server does not overwrite log information, and can store messages in non-volatile files (disks, for example).

The `enable syslog` command must be issued in order for messages to be sent to the remote syslog server(s). Syslog is disabled by default. A total of four syslog servers can be configured at one time.

When a syslog server is added, it is associated with the filter *DefaultFilter*. Use the `configure log target filter` command to associate a different filter.

The syslog facility level is defined as local0 – local7. The facility level is used to group syslog data.

Example

The following command configures the remote syslog server target with a critical severity:

```
configure syslog 123.45.67.78 local1 critical
```

History

This command was first available in ExtremeXOS 10.1.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure syslog delete

```
configure syslog delete [all | <ipaddress> | <ipPort>] {vr <vr_name>}
{local0 ... local7}
configure syslog delete <host name/ip> {: <udp-port>} [local0 ... local7]
```

Description

Deletes a remote syslog server address.

Syntax Description

all	Specifies all remote syslog servers.
ipaddress	Specifies the remote syslog server IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.

Default

If a virtual router is not specified, VR-Mgmt is used.

If a UDP port number is not specified, 514 is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command is used to delete a remote syslog server target.

Example

The following command deletes the remote syslog server with an IP address of 10.0.0.1:

```
configure syslog delete 10.0.0.1 local1
```

History

This command was first available in ExtremeXOS 10.1.

The ipPort parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create log filter

```
create log filter <name> {copy <filter name>}
```

Description

Creates a log filter with the specified name.

Syntax Description

name	Specifies the name of the filter to create.
copy	Specifies that the new filter is to be copied from an existing one.
filter name	Specifies the existing filter to copy.

Default

N/A.

Usage Guidelines

This command creates a filter with the name specified. A filter is a customizable list of events to include or exclude, and optional parameter values. The list of events can be configured by component or subcomponent with optional severity, or individual condition, each with optional parameter values. See the commands [configure log filter events](#) and [configure log filter events match](#) for details on how to add items to the filter.

The filter can be associated with one or more targets using the [configure log target filter](#) command to control the messages sent to those targets. The system has one built-in filter named *DefaultFilter*, which itself may be customized. Therefore, the [create log filter](#) command can be used if a filter other than *DefaultFilter* is desired. As its name implies, *DefaultFilter* initially contains the default level of logging in which every ExtremeXOS component and subcomponent has a pre-assigned severity level.

If another filter needs to be created that will be similar to an existing filter, use the `copy` option to populate the new filter with the configuration of the existing filter. If the `copy` option is not specified, the new filter will have no events configured and therefore no incidents will pass through it.

The total number of supported filters, including *DefaultFilter*, is 20.

Example

The following command creates the filter named *fdb2*, copying its configuration from the filter *DefaultFilter*:

```
create log filter fdb2 copy DefaultFilter
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete log filter

```
delete log filter [<filter name> | all]
```

Description

Deletes a log filter with the specified name.

Syntax Description

filter name	Specifies the filter to delete.
all	Specifies that all filters, except DefaultFilter, are to be deleted

Default

N/A.

Usage Guidelines

This command deletes the specified filter, or all filters except for the filter *DefaultFilter*. The specified filter must not be associated with a target. To remove that association, associate the target with *DefaultFilter* instead of the filter to be deleted, using the following command:

```
configure log target <target> filter DefaultFilter
```

Example

The following command deletes the filter named *fdb2*:

```
delete log filter fdb2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable cli-config-logging

```
disable cli-config-logging
```

Description

Disables the logging of CLI configuration commands to the switch Syslog.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Every command is displayed in the log window which allows you to view every command executed on the switch.

The `disable cli-config-logging` command discontinues the recording of all switch configuration changes and their sources that are made using the CLI via Telnet or the local console. After you disable configuration logging, no further changes are logged to the system log.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command disables the logging of CLI configuration command to the Syslog:

```
disable cli-config-logging
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable elsm ports

```
disable elsm ports <portlist>
```

Description

Disables the ELSM protocol for the specified ports.

Syntax Description

portlist	Specifies the port or ports for which ELSM should be disabled.

Default

The default is disabled.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages. To enable ELSM, use the following command:

```
enable elsm ports <portlist>
```

For more information about ELSM, see the command [enable elsm ports on page 733](#).

Example

The following command disables ELSM for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable elsm ports auto-restart

```
disable elsm ports <portlist> auto-restart
```

Description

Disable ELSM automatic restart for the specified ports.

Syntax Description

portlist	Specifies the port or ports for which ELSM auto-restart is being disabled.

Default

The default is enabled.

Usage Guidelines

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters and remains in the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and have the port enter the Down state by using one of the following commands:

```
clear elsm ports <portlist> auto-restart
enable elsm ports <portlist> auto-restart
```

If you use the `enable elsm ports <portlist> auto-restart` command, automatic restart is always enabled; you do not have to use the `clear elsm ports <portlist> auto-restart` command to clear the stuck state.

Enabling Automatic Restart. To enable ELSM automatic restart, you must explicitly configure this behavior on each ELSM-enabled port. If you enable ELSM automatic restart and an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

To enable automatic restart, use the following command:

```
enable elsm ports <portlist> auto-restart
```

If you configure automatic restart on one port, Extreme Networks recommends that you use the same configuration on its peer port.

Example

The following command disables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable log debug-mode

```
disable log debug-mode
```

Description

Disables debug mode. The switch stops logging events of severity debug-summary, debug-verbose, and debug-data.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to logging debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters
- Target format options process-name, process-id, source-function, and source-line

Example

The following command disables debug mode:

```
disable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable log display

```
disable log display
```

Description

Disables the sending of messages to the console display.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the log display is disabled, log information is no longer written to the serial console.

This command setting is saved to FLASH and determines the initial setting of the console display at boot up.

You can also use the following command to control logging to different targets:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Example

The following command disables the log display:

```
disable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable log target

```
disable log target [console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog [all |
<ipaddress> | <ipPort>] {vr <vr_name>} [local0 ... local7]]]
```

Description

Stops sending log messages to the specified target.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
	NOTE: This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
	NOTE: This parameter is available only on modular switches.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog host name or IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer, NVRAM, primary MSM, and backup MSM; all other targets are disabled by default.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command stops sending messages to the specified target. By default, the memory buffer, NVRAM, primary MSM, and backup MSM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the `session` target are in effect only for the duration of the console display or telnet session, and are not saved in FLASH. Changes to the other targets are saved to FLASH.

You can also use the following command to disable displaying the log on the console:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Modular Switches Only. Note that the `backup-msm` target is only active on the primary MSM, and the `primary-msm` target is only active on the backup MSM.

Example

The following command disables log messages to the current session:

```
disable log target session
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable rmon

```
disable rmon
```

Description

Disables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In a disabled state, the switch continues to respond queries of statistics. Collecting of history, alarms, and events is stopped; however, the switch still queries old data.

To view the status of RMON polling on the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | <memoryType>}
```

Example

The following command disables the collection of RMON statistics on the switch:

```
disable rmon
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable sflow

```
disable sflow
```

Description

Globally disables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables sFlow globally on the switch. When you disable sFlow globally, the individual ports are also put into the disabled state. If you later enable the global sFlow state, individual ports return to their previous state.

Example

The following command disables sFlow sampling globally:

```
disable sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

disable sflow ports

```
disable sflow ports <portlist>
```

Description

Disables sFlow statistical packet sampling and statistics gathering on a particular list of ports.

Syntax Description

portlist	Specifies a list of ports.

Default

Disabled.

Usage Guidelines

This command disables sFlow on a particular list of ports. Once sFlow is disabled on a port, sampling and polling will stop. If sFlow is disabled globally, all sampling and polling stops.

Use the following command to disable sFlow globally:

```
disable sflow
```

Example

The following command disables sFlow sampling on port 3:1:

```
disable sflow ports 3:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

disable sys-health-check

```
disable sys-health-check slot <slot>
```

Description

Discontinues sending backplane diagnostic packets on modular switches.

Syntax Description

slot	Specifies the slot to disable sending backplane diagnostic packets.

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when disabling backplane diagnostic packets, the following defaults apply:

- BlackDiamond 10808 and BlackDiamond 12804 switches—By default, the system health checker discontinues sending backplane diagnostic packets and returns the polling frequency to 60 seconds on the specified slot. Only polling is enabled.
- BlackDiamond 8800 series switch—By default, the system health checker discontinues sending backplane diagnostic packets to the specified slot. Only polling is enabled.

Usage Guidelines

When you use this command, backplane diagnostic packets are disabled and no longer sent by the system health checker.

BlackDiamond 10808 and BlackDiamond 12804 Switches Only. If you use this command, in addition to disabling backplane diagnostic packets, the polling frequency returns to the default of 60 seconds for the specified slot.

If you modify the `interval` in the `configure sys-health-check interval <interval>` command, the polling frequency is `interval*10` seconds. For example, if you configured an interval of 7, the polling frequency is now 70 seconds.

BlackDiamond 8800 Series Switch Only. If you modify the `interval` in the `configure sys-health-check interval <interval>` command and later disable backplane diagnostics, the configured interval for sending backplane diagnostic packets remains. The next time you enable backplane diagnostic packets, the health checker sends backplane diagnostics packets at the configured interval. For example, if you configure an interval of 8 seconds, the system health checker sends backplane diagnostic packets every 8 seconds.

To return to the "default" interval of 5 seconds, configure the frequency of sending backplane diagnostic packets to 5 seconds using the following command:

```
configure sys-health-check interval 5
```

Example

On the BlackDiamond 10808 and the BlackDiamond 12804 switches, the following example assumes that you did not modify the `interval` option in the `configure sys-health-check interval <interval>` command.

The following command disables backplane diagnostic packets on slot 3 and returns the polling frequency on slot 3 to 60 seconds:

```
disable sys-health-check slot 3
```

On the BlackDiamond 8800 series switch, the following example assumes that you did not modify the `interval` option in the `configure sys-health-check interval <interval>` command.

The following command disables backplane diagnostics on slot 3, polling is always enabled and occurs every 5 seconds.

```
disable sys-health-check slot 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

disable syslog

```
disable syslog
```

Description

Disables logging to all remote syslog server targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disables logging to all remote syslog server targets, not to the switch targets. This setting is saved in FLASH, and will be in effect upon boot up.

Example

The following command disables logging to all remote syslog server targets:

```
disable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable cli-config-logging

```
enable cli-config-logging
```

Description

Enables the logging of CLI configuration commands to the Syslog for auditing purposes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

ExtremeXOS allows you to record all configuration changes and their sources that are made using the CLI by way of Telnet or the local console. The changes are logged to the system log. Each log entry includes the user account name that performed the changes and the source IP address of the client (if Telnet was used). Configuration logging applies only to commands that result in a configuration change.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command enables the logging of CLI configuration commands to the Syslog:

```
enable cli-config-logging
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable elsm ports

```
enable elsm ports <portlist>
```

Description

Enables the ELSM protocol for the specified ports.

Syntax Description

portlist	Specifies the port or ports for which ELSM should be enabled.

Default

The default is disabled.

Usage Guidelines

The ELSM protocol allows you to detect CPU and remote link failures in the network. ELSM operates on a point-to-point basis; you only configure ELSM on the ports that connect to other devices within the network, but you must configure ELSM on both sides of the peer connections.

The Layer 2 connection between the ports determines the peer. You can have a direct connection between the peers or hubs that separate peer ports. In the first instance, the peers are also considered neighbors. In the second instance, the peer is not considered a neighbor.

An Extreme Networks device with ELSM enabled detects CPU and remote link failures by exchanging hello messages between two ELSM peers. If ELSM detects a failure, the ELSM-enabled port responds by blocking traffic on that port. For example, if a peer stops receiving messages from its peer, ELSM brings down that connection by blocking all incoming and outgoing data traffic on the port and notifying applications that the link is down.

Configuring and enabling ELSM on Extreme Networks devices running ExtremeXOS is backward compatible with Extreme Networks devices running ExtremeWare.

When you enable ELSM on a port, ELSM immediately blocks the port and it enters the Down state. When the port detects an ELSM-enabled peer, the peer ports exchange ELSM hello messages. At this point, the ports enter the transitional Down-Wait state. If the port receives Hello+ messages from its peer and does not detect a problem, the peers enter the Up state. If a peer detects a problem or there is no peer port configured, the port enters the Down state.

For more information about the types of ELSM hello messages, see the [configure elsm ports hellotime](#) command.



NOTE

ELSM and mirroring are mutually exclusive. You can enable either ELSM, or mirroring, but not both.

If you try to enable ELSM on a port that is already configured as a mirrored port, the switch displays messages similar to the following:

- Stand-alone switch

```
Cannot enable ELSM on port 1. Port is configured as mirror monitor port
```

- Modular switch

```
Cannot enable ELSM on port 1:1. Port is configured as mirror monitor port  
Configuration failed on backup MSM, command execution aborted!
```

Configuring the Hello Timer Interval. ELSM ports use hello messages to communicate information about the health of the network to its peer port. You can also configure the interval by which the ELSM-enabled ports sends hello messages. For more information about configuring the hello interval, see the command [configure elsm ports hellotime on page 657](#).

Disabling ELSM. ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states. To disable ELSM, use the following command:

```
disable elsm ports <portlist>
```

Example

The following command enables ELSM for slot 2, ports 1-2 on the switch:

```
enable elsm ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable elsm ports auto-restart

```
enable elsm ports <portlist> auto-restart
```

Description

Enables ELSM automatic restart for the specified ports.

Syntax Description

portlist	Specifies the port or ports for which ELSM auto-restart is being enabled.

Default

The default is enabled.

Usage Guidelines

You must explicitly configure this behavior on each ELSM-enabled port; this is not a global command.

By default, ELSM automatic restart is enabled. If an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports <portlist> auto-restart
enable elsm ports <portlist> auto-restart
```

If you use the `enable elsm ports <portlist> auto-restart` command, automatic restart is always enabled; you do not have to use the `clear elsm ports <portlist> auto-restart` command to clear the stuck state.

To disable automatic restart, use the following command:

```
disable elsm ports <portlist> auto-restart
```

If you configure automatic restart on one port, Extreme Networks recommends that you use the same configuration on its peer port.

Example

The following command enables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
enable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable log debug-mode

```
enable log debug-mode
```

Description

Enables debug mode. The switch allows debug events included in log filters to be logged.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables debug mode. Debug mode must be enabled prior to logging debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of `debug-summary`, `debug-verbose`, or `debug-data` when configuring filters
- Target format options `process-name`, `process-id`, `source-function`, and `source-line`

Example

The following command enables debug mode:

```
enable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable log display

```
enable log display
```

Description

Enables a running real-time display of log messages on the console display.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If you enable the log display on a terminal connected to the console port, your settings will remain in effect even after your console session is ended (unless you explicitly disable the log display).

You configure the messages displayed in the log using the `configure log display`, or `configure log target console-display` commands.

You can also use the following command to control logging to different targets:

```
enable log display
```

The `enable log display` command is equivalent to `enable log target console-display` command.

To change the log filter association, severity threshold, or match expression for messages sent to the console display, use the `configure log target console-display` command

Example

The following command enables a real-time display of log messages:

```
enable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable log target

```
enable log target [console | memory-buffer | nvram | primary-msm | primary-
node| backup-msm | backup-node| session | syslog [all | <ipaddress> |
<ipPort>] {vr <vr_name>} [local0 ... local7]]]
```

Description

Starts sending log messages to the specified target.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node of a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup node of a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer and NVRAM; all other targets are disabled by default.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command starts sending messages to the specified target. By default, the memory-buffer, NVRAM, primary MSM, and backup MSM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the `session` target are in effect only for the duration of the console display or Telnet session, and are not saved in FLASH. Others are saved in FLASH.

You can also use the following command to enable displaying the log on the console:

```
enable log display
```

The `enable log display` command is equivalent to the `enable log target console-display` command.

Modular Switches Only. Note that the `backup-msm` target is only active on the primary MSM, and the `primary-msm` target is only active on the backup MSM.

Example

The following command enables log messages on the current session:

```
enable log target session
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable rmon

```
enable rmon
```

Description

Enables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In an enabled state, the switch responds to the following four groups:

- **Statistics**—The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.
- **History**—The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The group features user-defined sample intervals and bucket counters for complete customization of trend analysis.
- **Alarms**—The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. In addition, alarm thresholds may be auto calibrated or set manually.
- **Events**—The Events group creates entries in an event log and/or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. The action taken can be configured to ignore it, to log the event, to send an SNMP trap to the receivers listed in the trap receiver table, or to both log and send a trap. The RMON traps are defined in RFC 1757 for rising and falling thresholds.

The switch also supports the following parameters for configuring the RMON agent, as defined in RFC 2021:

- **probeCapabilities**—If you configure the probeCapabilities object, you can view the RMON MIB groups supported on at least one interface by the probe.
- **probeSoftwareRev**—If you configure the probeSoftwareRev object, you can view the current software version of the monitored device.
- **probeHardwareRev**—If you configure the probeHardwareRev object, you can view the current hardware version of the monitored device.

- **probeDateTime**—If you configure the **probeDateTime** object, you can view the current date and time of the probe.
- **probeResetControl**—If you configure the **probeResetControl** object, you can restart a managed device that is not running normally. Depending on your configuration, you can do one of the following:
 - **Warm boot**—A warm boot restarts the device using the current configuration saved in non-volatile memory.
 - **Cold boot**—A cold boot causes the device to reset the configuration parameters stored in non-volatile memory to the factory defaults and then restarts the device using the restored factory default configuration.

**NOTE**

You can only use the RMON features of the system if you have an RMON management application and have enabled RMON on the switch.

RMON requires one probe per LAN segment, and stand-alone RMON probes have traditionally been expensive. Therefore, the approach taken by Extreme Networks has been to build an inexpensive RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management. The switch accurately maintains RMON statistics at the maximum line rate of all of its ports.

For example, statistics can be related to individual ports. Also, because a probe must be able to see all traffic, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the switch means that all ports can have security features enabled.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | <memoryType>}
```

Example

The following command enables the collection of RMON statistics on the switch:

```
enable rmon
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable sflow

```
enable sflow
```

Description

Globally enables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables sFlow globally on the switch.



NOTE

sFlow and mirroring are mutually exclusive on the BlackDiamond 8800 original modules and the Summit X450 series switch (whether or not included in a SummitStack). You can enable either sFlow, or mirroring on these modules, but not both. This restriction is not applicable to the BlackDiamond 10808 or the BlackDiamond 12804 switches.

sFlow and mirroring are not mutually exclusive on the BlackDiamond 8800 a-series and e-series modules and the Summit X450a and X450e series switches (whether or not included in a SummitStack). You can enable sFlow and mirroring at the same time.

If you have a combination of BlackDiamond 8800 original, a-series, and e-series modules installed, or if you have a combination of X450 and X450a or X450e switches connected in a SummitStack, with sFlow configured globally on the switch or stack, you can enable mirroring on original I/O module ports or X450 ports provided you do not enable sFlow on original I/O module ports or X450 ports.

Beginning with ExtremeXOS 11.5, you can enable sFlow and mirroring at the same time on the following platforms:

- BlackDiamond 8800 e-series modules
- BlackDiamond 8800 a-series modules
- Summit X450e-series switches
- Summit X450a-series switches

BlackDiamond 8800 Original Modules and Summit X450 Series Switch Only. sFlow uses QP2 to sample traffic on the BlackDiamond 8800 original modules and the Summit X450 series switch. If you enable sFlow globally on those platforms, the switch displays a message similar to the following:

SFLOW will use Qos Profile QP2.

Any traffic grouping using QP2 may encounter unexpected results when sFlow is enabled. For more information about QoS, see [Chapter 14, “Quality of Service,”](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command enables sFlow sampling globally:

```
enable sflow
```

If you enable sFlow globally on the Summit X450 series switch or the BlackDiamond 8800 original modules, the switch displays a message similar to the following:

```
SFLOW will use Qos Profile QP2.
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

enable sflow ports

```
enable sflow ports <port_list>
```

Description

Enables sFlow statistical packet sampling on a particular list of ports.

Syntax Description

port_list	Specifies a list of ports.

Default

Disabled.

Usage Guidelines

This command enables sFlow on a particular list of ports. You also need to enable sFlow globally in order to gather statistics and send the data to the collector. Once sFlow is enabled globally, and on the ports of interest, sampling and polling will begin.

Use the following command to enable sFlow globally:

```
enable sflow
```



NOTE

sFlow and mirroring are mutually exclusive on the BlackDiamond 8800 original modules and the Summit X450 series switch (whether or not included in a SummitStack). You can enable either sFlow, or mirroring on these modules, but not both. This restriction is not applicable to the BlackDiamond 10808 or the BlackDiamond 12804 switches.

sFlow and mirroring are not mutually exclusive on the BlackDiamond 8800 a-series and e-series modules and the Summit X450a and X450e series switches (whether or not included in a SummitStack). You can enable sFlow and mirroring at the same time.

If you have a combination of BlackDiamond 8800 original, a-series, and e-series modules installed, or if you have a combination of X450 and X450a or X450e switches connected in a SummitStack, with sFlow configured globally on the switch or stack, you can enable mirroring on original I/O module ports or X450 ports provided you do not enable sFlow on original I/O module ports or X450 ports.



NOTE

Beginning with ExtremeXOS 11.5, you can enable sFlow and mirroring at the same time on the following platforms:

- BlackDiamond 8800 e-series modules
- BlackDiamond 8800 a-series modules

- Summit X450e-series switches
- Summit X450a-series switches

BlackDiamond 8800 Series Switch and SummitStack Only. On the BlackDiamond 8800 series switch with a combination of a-series, e-series, and original modules installed, or on a SummitStack that contains X450 and X450a, X450e, or X250e switches, please note the following sFlow restrictions:

- If you enable sFlow globally on the switch and on original I/O module ports or on X450 ports and you try to enable mirroring on original, a-series, or e-series I/O module ports or x450 ports, the switch displays a message similar to the following:

```
Error: Mirroring is not compatible with SFlow. Mirroring not enabled!
```

```
Configuration failed on backup MSM, command execution aborted!
```

If you see the above message, you must disable sFlow on all of the original I/O module ports or X450 ports.

- If you enable sFlow globally on the switch and on a-series or e-series I/O module ports or on X450 ports and you try to enable mirroring on original I/O module ports or X450 ports, the switch does not display an error. You can successfully enable mirroring on original I/O module ports or X450 ports, provided you do not enable sFlow on original I/O module ports or X450 ports.
- If you enable sFlow on original I/O module ports or X450 ports, sFlow uses QP2 to sample traffic. When enabling sFlow on those ports, the switch displays a message similar to the following:

```
SFLOW will use Qos Profile QP2.
```

For more information about the BlackDiamond 8800 a-series, e-series modules, and original modules, and X450a, X450e, X250e and X450 switches, see [Chapter 2, “ExtremeXOS Overview”](#) in the *ExtremeXOS Concepts Guide* and the *Extreme Networks Consolidated ExtremeXOS Hardware Guide*. For more information about mirroring, see [Chapter 6, “Configuring Slots and Ports on a Switch”](#) in the *ExtremeXOS Concepts Guide*. For more information about QOS, see [Chapter 14, “Quality of Service”](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command enables sFlow sampling on the port 3:1:

```
enable sflow ports 3:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

enable sys-health-check

```
enable sys-health-check slot <slot>
```

Description

Enables backplane diagnostic packets on the specified slot.

Syntax Description

slot	Specifies the slot to participate in sending backplane diagnostic packets.

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when enabling backplane diagnostic packets, the following defaults apply:

- BlackDiamond 10808 and BlackDiamond 12804 switches—The system health checker tests the packet path every 6 seconds for the specified slot.
- BlackDiamond 8800 series switch—The system health checker tests the data link every 5 seconds for the specified slot.

Usage Guidelines

Configure the system health checker with guidance from Extreme Networks Technical Support personnel.

The system health checker tests I/O modules and the backplane by sending diagnostic packets. On the BlackDiamond 10808 and the BlackDiamond 12804 switches, additional checking for the validity of these packets is completed by performing a checksum. By isolating faults to a specific module or backplane connection, the system health checker notifies you of a possible hardware failure.

System health check errors are reported to the syslog. Syslog output includes the slot number where the problem occurred, the loopback packet ID number, and a notification that the MSM did not receive the last packet. If you see an error, please contact Extreme Networks Technical Support.



NOTE

Enabling backplane diagnostic packets increases CPU utilization and competes with network traffic for resources.

The system health checker continues to periodically forward test packets to failed components.

To configure the frequency of the backplane diagnostic packets, use the `configure sys-health-check interval` command.

BlackDiamond 10808 and BlackDiamond 12804 Switches Only. If you enable backplane diagnostic packets on a slot, the polling timer changes from its current value to the backplane diagnostic packet interval configured on that slot. For example, if the frequency of sending backplane diagnostic packets is 7 seconds, the polling value is 7 seconds.

Displaying the System Health Check Setting. To display the system health check polling setting on the switch, use the following command:

```
show switch
```

As previously described, polling is always enabled on the switch, which is why you see the system health check setting as Enabled. The following truncated output from a BlackDiamond 8810 switch displays the system health check setting (displayed as SysHealth check):

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A2:60

SysHealth check:  Enabled
Recovery Mode:    None
System Watchdog:  Enabled
```

Example

The following command enables backplane diagnostic packets on slot 6:

```
enable sys-health-check slot 6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

enable syslog

```
enable syslog
```

Description

Enables logging to all remote syslog host targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable remote logging, you must do the following:

- Configure the syslog host to accept and log messages.
- Enable remote logging by using the `enable syslog` command.
- Configure remote logging by using the `configure syslog` command.

When you use the `enable syslog` command, the exporting process of the syslog begins. This command also determines the initial state of an added remote syslog target.

Example

The following command enables logging to all remote syslog hosts:

```
enable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show elsm

```
show elsm
```

Description

Displays summary information for all of the ELSM-enabled ports on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the operational state of ELSM on the configured ports.

If no ports are configured for ELSM, the switch does not display any information.

For ELSM-enabled ports, this command displays the following information in a tabular format:

- Port—The port number of the ELSM-enabled port.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer.
If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state.
 - Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.
If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.
 - Down-Wait—Indicates a transitional state.
If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hello threshold (by default 2 messages), the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.
 - Down-Stuck—Indicates that the port is down and requires user intervention.
If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:

By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` on page 735 for more information.

If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports <portlist> auto-restart
enable elsm ports <portlist> auto-restart
```

- Hello time —The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.

Additional show Command. You can also use the `show ports {<port_list>} information {detail}` command to display ELSM information.

If you do not specify the `detail` parameter, the following columns display the current state of ELSM on the switch:

- Flags
 - L—Indicates that ELSM is enabled on the switch
 - - —Indicates that ELSM is disabled on the switch
- ELSM
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional `detail` parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch)
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM
 - Enabled—Indicates that ELSM is enabled on the switch
 - Disabled—Indicates that ELSM is disabled on the switch

Example

The following command displays summary configuration information for all of the ELSM-enabled ports on the switch:

```
show elsm
```

The following is sample output from this command:

Port	ELSM State	Hello Time
5:14	Up	1 (second)
5:18	Down	1 (second)

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show elsm ports

```
show elsm ports <all | portlist>
```

Description

Displays detailed information for one or more ELSM-enabled ports.

Syntax Description

all	Displays detailed ELSM information for all ports.
portlist	Displays detailed ELSM information for one or more ports.

Default

N/A.

Usage Guidelines

Use this command to display detailed information about the operational state of ELSM on the configured ports.

This command displays in a tabular format the following ELSM data for one or more ELSM-enabled ports on the switch:

- Port—The port number of the ELSM-enabled port.
- Link State—The state of the link between ELSM-enabled (peer) ports. The link state can be one of the following:
 - Ready—Indicates that the port is enabled but there is no link.
 - Active—Indicates that the port is enabled and the physical link is up.
- ELSM Link State—The current state of the ELSM logical link on the switch. The ELSM link state can be one of the following:
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer. If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state. The up timer is 6* hello timer, which by default is 6 seconds.

- Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.
If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.
- Down-Wait—Indicates a transitional state.
If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hold threshold, the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.
- Down-Stuck—Indicates that the port is down and requires user intervention.
If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:
By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` on page 735 for more information.
If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports <portlist> auto-restart
enable elsm ports <portlist> auto-restart
```
- Hello Transmit State—The current state of ELSM hello messages being transmitted. The transmit state can be one of the following:
 - HelloRx(+)—Specifies that the ELSM-enabled port is up and receiving Hello+ messages from its peer. The port remains in the HelloRx+ state and restarts the HelloRx timer each time it receives a Hello+ message. If the HelloRx timer expires, the hello transmit state enters HelloRX(-). The HelloRx timer is 6 * hello timer, which by default is 6 seconds.
 - HelloRx(-)—Specifies that the ELSM-enabled port either transitions from the initial ELSM state or is up and not receiving hello messages because there is a problem with the link or the peer is missing.
- Hello time—The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.
- Hold Threshold—The number of Hello+ messages required by the ELSM-enabled port to transition from the Down-Wait state to the Up state within the hold threshold.
- UpTimer Threshold—The number of hello times that span without receiving Hello+ packets before a port changes its ELSM state from Up to Down.
- Auto Restart—The current state of ELSM automatic restart on the port. The state of Auto Restart can be one of the following:
 - Enabled—If an ELSM-enabled port goes down, ELSM automatically brings up the down port. This is the default behavior.
 - Disabled If an ELSM-enabled port goes down, the port enters and remains in the Down-Stuck state until you clear the stuck state.

For more information about automatic restart, see the command `enable elsm ports auto-restart` on page 735.
- Sticky Threshold—Specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

- Sticky Threshold Counter—The number of times the port transitions from the Up state to the Down state.
- Down Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Down to Up. When ELSM is enabled on a port and it is in a Down state, before it changes its ELSM state from Down to Up, it expects to receive at least a “Hold Threshold” number of Hello+ packets during the Down Timeout period after it receives the first Hello+ packet from its peer. It is equal to $[\text{Hello Time} * (\text{Hold Threshold} + 2)]$.
- Up Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Up to Down after receiving the last Hello+ packets. When a port is in an Up state, it expects to receive a Hello+ packet from its peer every “Hello Time” period to maintain its Up state. When it does not receive a Hello+ packet after an “Up Timeout” period, it changes its ELSM state from Up to Down. It is equal to $[\text{Hello Time} * \text{UpTimer Threshold}]$.

The remaining output displays counter information. Use the counter information to determine the health of the ELSM peers and how often ELSM has gone up or down. The counters are cumulative.

- RX Hello+—The number of Hello+ messages received by the port.
- Rx Hello- —The number of Hello- messages received by the port.
- Tx Hello+—The number of Hello+ messages sent by the port.
- Tx Hello- —The number of Hello- messages sent by the port.
- ELSM Up/Down Count—The number of times ELSM has been up or down.

To clear, reset the counters, use either the `clear elsm {ports <portlist>} counters` or the `clear counters` command.

Additional show Command. You can also use the `show ports {<port_list>} information {detail}` command to display ELSM information.

If you do not specify the `detail` parameter, the following columns display the current state of ELSM on the switch:

- Flags
 - L—Indicates that ELSM is enabled on the switch
 - - —Indicates that ELSM is disabled on the switch
- ELSM
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional `detail` parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch)
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.

- Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM
 - Enabled—Indicates that ELSM is enabled on the switch
 - Disabled—Indicates that ELSM is disabled on the switch

Example

The following command displays detailed ELSM information for all configured ports on the switch:

```
show elsm ports all
```

The following is sample output from this command:

```
ELSM Info Port 4:4
  Link State           : Active
  ELSM Link State      : Up
  ELSM State           : Up
  Hello Transmit State  : HelloRx(+)
  Hello Time           : 100 msec
  Hold Threshold       : 2
  UpTimer Threshold    : 6
  Auto Restart         : Enabled
  Down Timeout         : 400 msec
  Up Timeout           : 600 msec
  Rx Hello+            : 667960
  Rx Hello-            : 0
  Tx Hello+            : 667958
  Tx Hello-            : 0
  ELSM Up/Down Count   : UP: 0    DOWN: 0

ELSM Info Port 4:4
  Link State           : Active
  ELSM Link State      : Up
  ELSM State           : Up
  Hello Transmit State  : HelloRx(+)
  Hello Time           : 100 msec
  Hold Threshold       : 2
  UpTimer Threshold    : 6
  Auto Restart         : Disabled
  Sticky Threshold     : 1
  Sticky Threshold Counter : 0
  Down Timeout         : 400 msec
  Up Timeout           : 600 msec
  Rx Hello+            : 708204
  Rx Hello-            : 0
  Tx Hello+            : 708201
  Tx Hello-            : 0
  ELSM Up/Down Count   : UP: 0    DOWN: 0
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show fans

```
show fans {detail}
```

Description

Displays the status of the fans in the system.

Syntax Description

detail	The detail option is reserved for future use.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the fans.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects and displays the following fan information:

- State—The current state of the fan. Options are:
 - Empty: There is no fan installed.
 - Failed: The fan failed.
 - Operational: The fan is installed and working normally.
- NumFan—The number of fans in the fan tray.
- Fan Name, displayed as Fan-1, Fan-2, and so on (modular switches also include a description of the location, for example, Upper or Upper-Right)—Specifies the individual state for each fan in a fan tray and its current speed in revolutions per minute (rpm).

On modular switches, the output also includes the following information:

- PartInfo—Information about the fan tray, including the:
 - Serial number—A collection of numbers and letters, that make up the serial number of the fan. This is the first series of numbers and letters in the display.
 - Part number—A collection of numbers and letters, that make up the part number of the fan. This is the second series of numbers and letters in the display.
- Revision—The revision number of the fan.
- Odometer—Specifies the power-on date and how long the fan tray has been operating since it was first powered-on.
- Temperature—Specifies, in celsius, the current temperature of the fan. (BlackDiamond 10808 switch only.)

Example

The following command displays the status of the installed fans. If a fan is not installed, the state of the fan is `Empty`.

```
show fans
```

The following is sample output from a BlackDiamond 10808 switch:

Right(Rear-facing) FanTray 1 information:

```
State:           Operational
NumFan:          2
PartInfo:        0340F-00046 454301-00-03
Revision:        3.0
Odometer:        555 days 4 hours 45 minutes 16 seconds since Mar-19-2004
Temperature:     34.25 deg C
Upper Fan 1:     Operational at 4380 rpms
Lower Fan 2:     Operational at 4380 rpms
```

Left(Rear-facing) FanTray 2 information:

```
State:           Operational
NumFan:          2
PartInfo:        0340F-00047 454301-00.03
Odometer:        11845 days 6 hours 55 minutes 39 seconds since Mar-08-2004
Temperature:     28.93 deg C
Upper Fan 1:     Operational at 4380 rpms
Lower Fan 2:     Operational at 4320 rpms
```

The following is sample output from the BlackDiamond 8800 series switch:

FanTray information:

```
State:           Operational
NumFan:          9
PartInfo:        0404X-00015 450102-00-01
Revision:        1.0
Odometer:        111 days 16 hours 30 minutes   since Oct-13-2004
Upper-Left  Fan-1: Operational at 2880 RPM
Middle-Left Fan-2: Operational at 2820 RPM
Lower-Left  Fan-3: Operational at 2820 RPM
Upper-Center Fan-4: Operational at 2820 RPM
Center      Fan-5: Operational at 2820 RPM
Lower-Center Fan-6: Operational at 2880 RPM
Upper-Right Fan-7: Operational at 2880 RPM
Middle-Right Fan-8: Operational at 2820 RPM
Lower-Right Fan-9: Operational at 2880 RPM
```

The following is sample output from the BlackDiamond 12804 switch:

FanTray information:

```
State:           Operational
NumFan:          6
PartInfo:        05304-00056 454030-00-03
Revision:        3.0
Odometer:        66 days 10 hours 30 minutes   since Jul-26-2005
Upper-Left  Fan-1: Operational at 2040 RPM
Lower-Left  Fan-2: Operational at 2040 RPM
Upper-Center Fan-3: Operational at 2040 RPM
```

```

Lower-Center Fan-4:      Operational at 2040 RPM
Upper-Right  Fan-5:      Operational at 2040 RPM
Lower-Right  Fan-6:      Operational at 2040 RPM

```

The following is sample output from a SummitX450 switch:

```

FanTray information:
  State:                Operational
  NumFan:               1
  Fan-1:                Operational at 1000 RPM

```

The following is a sample output from a SummitStack:

```

FanTray-1 information:
  State:                Operational
  NumFan:               1
  Fan-1:                Operational at 1000 RPM

FanTray-2 information:
  State:                Operational
  NumFan:               1
  Fan-1:                Operational at 1000 RPM

FanTray-3 information:
  State:                Operational
  NumFan:               1
  Fan-1:                Operational at 1000 RPM

FanTray-4 information:
  State:                Operational
  NumFan:               1
  Fan-1:                Operational at 1000 RPM

FanTray-5 information:
  State:                Empty

FanTray-6 information:
  State:                Empty

FanTray-7 information:
  State:                Empty

FanTray-8 information:
  State:                Empty

```

History

This command was first available in an ExtremeXOS 10.1.

Information about the location of the fan tray for the BlackDiamond 10808 switch (right or left fan tray) was added to the `show fans` output in ExtremeXOS 11.0.

Information about the location of the fan tray for the BlackDiamond 8810 switch (upper-left, middle left, lower-left, upper- lower-center, upper-right, middle-right, and lower-right) was added to the `show fans` output in ExtremeXOS 11.1.

Information about the current speed in rpm for the Summit family of switches was added to the `show fans` output in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show log

```
show log {messages [memory-buffer | nvram]} {events {<event-condition> |
<event-component>}} {<severity> {only}} {starting [date <date> time <time>
| date <date> | time <time>]} {ending [date <date> time <time> | date
<date> | time <time>]} {match <regex>} {chronological}
```

Description

Displays the current log messages.

Syntax Description

messages	Specifies the target location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory (default).
nvram	Show messages stored in NVRAM.
events	Show event messages.
event-condition	Specifies the event condition to display.
event-component	Specifies the event component to display.
severity	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed
starting	Show messages with timestamps equal to or greater than that specified
date	Specifies the date, where date is <month (1-12)> / <day (1-31)> [/ <year (yyyy)>].
time	Specifies the time, where time is <hour (0-23)> {: <minute (0-59)> {: <seconds (0-59)> {. <hundredths>}}}
ending	Show messages with timestamps equal to or less than that specified.
regex	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies displaying log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer
- event—no restriction (displays user-specified event)
- severity—none (displays everything stored in the target)
- starting, ending—if not specified, no timestamp restriction
- match—no restriction
- chronological—if not specified, show messages in order from newest to oldest

Usage Guidelines

Switch configuration and fault information is filtered and saved to target logs, in a memory buffer, and in NVRAM. Each entry in the log contains the following information:

- **Timestamp**—records the month and day of the event, along with the time (hours, minutes, seconds, and hundredths).
- **Severity Level**—indicates the urgency of a condition reported in the log. [Table 17](#) describes the severity levels assigned to events.
- **Component, Subcomponent, and Condition Name**—describes the subsystem in the software that generates the event. This provides a good indication of where a fault might lie.
- **Message**—a description of the event occurrence. If the event was caused by a user, the user name is also provided.

This command displays the messages stored in either the internal memory buffer or in NVRAM. The messages shown can be limited by specifying a severity level, a time range, or a match expression. Messages stored in the target have already been filtered as events occurred, and specifying a severity or match expression on the `show log` command can only further limit the messages shown.

If the `messages` keyword is not present, the messages stored in the memory-buffer target are displayed. Otherwise, the messages stored in the specified target are displayed.

If the `only` keyword is present following the severity value, then only the events at that exact severity are included. Without the `only` keyword, events at that severity or more urgent are displayed. For example, `severity warning` implies critical, error, or warning, whereas `severity warning only` implies only warning.

Messages whose timestamps are equal or later than the starting time and are equal or earlier than the specified ending time will be shown if they also pass the severity requirements and match expression, if specified.

If a `match` phrase is specified, the formatted message must match the simple regular expression specified by `match-expression` for it to be shown.

A simple regular expression is a string of single characters including the dot character (`.`), which are optionally combined with quantifiers and constraints. A dot matches any single character while other characters match only themselves (case is significant). Quantifiers include the star character (`*`) that matches zero or more occurrences of the immediately preceding character or dot. Constraints include the caret character (`^`) that matches at the beginning of a message, and the currency character (`$`) that matches at the end of a message. Bracket expressions are not supported. There are a number of sources available on the Internet and in various language references describing the operation of regular expressions.

If the `chronological` keyword is specified, messages are shown from oldest to newest; otherwise, messages are displayed newest to oldest.

Severity Level. The severity levels are `critical`, `error`, `warning`, `notice`, and `info`, plus three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`. In log messages, the severity levels are shown by four letter abbreviations. The abbreviated forms are:

- Critical—Crit
- Error—Erro
- Warning—Warn
- Notice—Noti

- Info—Info
- Debug-Summary—Summ
- Debug-Verbose—Verb
- Debug-Data—Data

The three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`, require that debug mode be enabled (which may cause a performance degradation). See the command `enable log debug-mode` on page 737. Table 17 describes the security levels.

Table 17: Severity levels assigned by the switch

Level	Description
Critical	A serious problem has been detected that is compromising the operation of the system and that the system cannot function as expected unless the situation is remedied. The switch may need to be reset.
Error	A problem has been detected that is interfering with the normal operation of the system and that the system is not functioning as expected.
Warning	An abnormal condition, not interfering with the normal operation of the system, has been detected that may indicate that the system or the network in general may not be functioning as expected.
Notice	A normal but significant condition has been detected, which signals that the system is functioning as expected.
Info (Informational)	A normal but potentially interesting condition has been detected, which signals that the system is functioning as expected and simply provides information or confirmation about the condition.
Debug-Summary	A condition has been detected that may interest a developer determining the reason underlying some system behavior.
Debug-Verbose	A condition has been detected that may interest a developer analyzing some system behavior at a more verbose level than provided by the debug summary information.
Debug-Data	A condition has been detected that may interest a developer inspecting the data underlying some system behavior.

Log entries remain in the NVRAM log after a switch reboot. Issuing a `clear log` command does not remove these static entries. To remove log entries from NVRAM, use the following command:

```
clear log messages nvram
```

Example

The following command displays messages with a critical severity:

```
show log critical
```

The following command displays messages with warning, error, or critical severity:

```
show log warning
```

The following is sample output from a modular switch:

```
11/12/2004 00:38:10.30 <Warn:dm.Warn> MSM-A: Insufficient Power to power-on Slot-7

11/12/2004 00:38:08.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insufficient power
```

```
11/12/2004 00:36:23.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insufficient power
```

```
...
```

A total of 83 log messages were displayed.

The following command displays messages containing the string “slot 2”:

```
show log match "slot 2"
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log components

```
show log components {<event component>} {version}
```

Description

Displays the name, description and default severity for all components.

Syntax Description

event component	Specifies the component to display.
version	Specifies the version number of the component.

Default

N/A.

Usage Guidelines

This command displays the name, description, and default severity defined for the specified components or subcomponents.

Depending on the software version running on your switch or your switch model, additional or different component information might be displayed.

Example

The following command displays the log components:

```
show log components
```

The following is sample output from this command:

Severity			
Component		Title	Threshold

AAA		Authentication, Authorization, Accounting	Info
	RADIUS	Remote Authentication Dial In User Service	Error
	TACACS	Terminal Access Controller Access Control Syst	Info
ACL		ACL	Info
	CLEARFlow	CLEARFlow	Info
	Policy	Policy actions	Info
bgp		Border Gateway Protocol	Info
	damp	BGP Route Flap Dampening related debug message	Error
	event	BGP FSM related events	Error
	inUpdt	Incoming Update related debug msgs	Warning
	keepalive	BGP keepalive message	Warning
	misc	Miscellaneous debug (Import, Aggregate, NextHop	Warning
	msgs	Debug for BGP messages (OPEN, Update, Notifica	Warning
	outUpdt	Transmit Update related debug	Warning

bootp		BOOTP, DHCP Component	Error
	relay	BOOTP Relay trace component	Error
	server	DHCP Server subcomponent	Info
cli		Command Line Interface	Info
	shell	CLI configuration shell.	Error
	subagent	CLI application subagent	Error
cm		Configuration Manager	Warning
	file	CM file operation events	Warning
	sys	CM system events	Warning
DM		Device Manager	Info
	Card	Device Manager Card State Machine	Info
dosprot		dosprot	Info
ds		Directory Services	Error
EAPS		Ethernet Automatic Protection Switching	Info
	SharedPort	EAPS SharedPort Domain	Info
EDP		Extreme DIScovery Protocol (EDP)	Error
ELRP			
	Report	Extreme Loop Recognition Protocol	Warning
EPM		Extreme Process Manager	Info
	KLM	Kernel Loadable Module Manager	Notice
	Msg	Message Handler	Info
	Upgrade	Upgrade Manager	Info
	Version	Version Manager	Critical
ESRP		Extreme Standby Router Protocol	Error
	Aware	Subsystem description	Info
	InPdu	Subsystem description	Info
	Nbr	Subsystem description	Info
	OutPdu	Subsystem description	Info
	State	ESRP State Transitions	Warning
	System	Subsystem description	Warning
	Track	Subsystem description	Warning
	Vlan	Extreme Standby Router Protocol	Info
fdb		fdb module event	Error
HAL		Hardware Abstraction Layer	Error
	Card	Card State Driver	Info
	FDB	Forwarding Database Driver	Info
	IPv4ACL	IPv4 Access Control List Driver	Info
	IPv4Adj	IPv4 Adjacency Driver	Info
	IPv4FIB	IPv4 FIB Driver	Info
	IPv4Mc	IPv4 Multicast Driver	Info
	Mirror	Mirroring Driver	Error
	Msg	Message Handler	Info
	Port	I/O Port Driver	Info
	SM	Switch Manager	Info
	Sys	System Driver	Info
	VLAN	VLAN Driver	Info
IPMC		IP Multicast Main Module	Info
	Snoop	IP Multicast Snooping Module	Error
	VLAN	IP Multicast VLAN Module	Error
Kern		Kernel messages	Error
LACP		Link Aggregation Control Protocol	Info
lldp		Link Layer Discovery Protocol (IEEE 802.1AB)	Warning
log		Log server messages	Warning
netTool		netTools framework	Error
	dnsclient	Dns Client	Error
	dnsproxy	Dns Proxy	Error

	routeradv	IPv6 Router Advertisements	Warning
	sntp	Sntp client	Warning
nl		Network Login	Info
	dot1x	802.1x-based Network Login	Warning
	mac	MAC-based Network Login	Warning
	web	Web-based Network Login	Warning
NM		Node Manager	Info
ospf		open shortest path first	Error
	event	ospf events	Info
	hello	ospf hello	Error
	lsa	ospf link-state advertisement	Error
	neighbor	ospf neighbor	Error
	spf	ospf shortest path first	Error
ospfv3		OSPFv3 related EMS messages	Warning
	events	OSPF6 events related messages	Error
	lsa	LSA related messages	Warning
	nbr	OSPF6 neighbor related EMS messages	Warning
	pkt	OSPF6 Packet receive/transmit/processing relat	Warning
	route	OSPF6 route add/delete related messages	Warning
	spf	SPF computation related messages	Error
pim		Pim Protocol Events	Warning
	cache	PIM cache maintenance.	Warning
	debug	PIM debug messages	Notice
	hello	Hello messages	Warning
	mcdbg	multicast forwarding engine	Warning
	msg	Trace for pim control packtes	Notice
	nbr	Neighbor creation/deletion etc	Warning
	rpm	RP message exchange.	Warning
pm		Policy Manager	Error
	config	Policy file events	Info
POE		Inline Power	Notice
rip		RIP routing	Error
	cfg	rip configuration	Warning
	event	rip events	Warning
	inUpdt	rip - inbound route updates	Warning
	msgs	rip - socket messages in and out	Warning
	outUpdt	rip - outbound route updates	Warning
	sys	rip - exos kernel interface	Warning
ripng		RIPng Protocol Events	Warning
	debug	RIPng debug messages	Notice
	external	RIPng external interface related messages	Warning
	message	RIPng control messages	Warning
	route	Hello messages	Warning
rmon		RMON general info	Error
	alarm	RMON alarm info	Error
	estat	RMON statistics info	Error
	event	RMON event info	Error
	history	RMON history	Error
RtMgr		Route Manager	Info
	VLAN	rtmgr vlan interface	Info
sflow		Sflow Protocol Events	Warning
	debug	SFLOW debug messages	Notice
	extended	SFLOW extended data collection	Notice
	msg	SFLOW process initializaion related message	Warning
	sample	SFLOW sample collection related messages	Warning
	statistics	SFLOW port statistics related message	Warning

STP		Spanning-Tree Protocol	Error
	InBPDU	STP In Bridge Protocol Data Unit	Warning
	OutBPDU	STP Out Bridge Protocol Data Unit	Warning
	System	STP System	Error
System		XOS system related log messages	Info
telnetd		telnet server	Info
tftpd		tftp server	Info
thttpd		thttp server	Info
trace		Debug trace messages	Warning
vlan		Vlan mgr	Info
	ack	vlan ack	Error
	dbg	Debug information	Info
	err	errors	Error
	mac	Virtual MAC Debugging	Info
	msgs	Messages	Info
VRRP		Config/State messages	Warning
	Advert	Subsystem description	Warning
	System	System/Library messages	Warning

A total of 143 component(s) were displayed.

The following command displays the version number of the VRRP component:

```
show log components vrrp version
```

The following is sample output from this command:

Component	Title	Version
-----	-----	-----
VRRP	Config/State messages	2.4
Advert	Subsystem description	3.1
System	System/Library messages	3.2

A total of 3 component(s) were displayed.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log configuration

```
show log configuration
```

Description

Displays the log configuration for switch log settings, and for certain targets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the log configuration for all targets. The state of the target, enabled or disabled is displayed. For the enabled targets, the associated filter, severity, match expression, and format is displayed. The debug mode state of the switch is also displayed.

Example

The following command displays the configuration of all the log targets and all existing filters:

```
show log configuration
```

The following is sample output from this command:

```
Debug-Mode: Enabled
```

```
Log Target      : memory-buffer
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Debug-Data (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
  Buffer size  : 1000 messages

Log Target      : nvram
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
```

```

Log Target      : console
  Enabled ?    : no
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Info (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

```

Log Filter Name: DefaultFilter
I/
E  Comp.    Sub-comp.    Condition    Severity
-  -----  -
I  All

```

```

Log Filter Name: myFilter
I/
E  Comp.    Sub-comp.    Condition    Severity
-  -----  -
I  STP

```

Include/Exclude: I - Include, E - Exclude
 Component Unreg: * - Component/Subcomponent is not currently registered
 Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
 Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
 + - Debug Severities, but log debug-mode not enabled
 If Match parameters present:
 Parameter Flags: S - Source, D - Destination, (as applicable)
 I - Ingress, E - Egress, B - BGP
 Parameter Types: Port - Physical Port list, Slot - Physical Slot #
 MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
 VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
 L4 - Layer-4 Port #, Num - Number, Str - String
 Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
 Proc - Process Name
 Strict Match : Y - every match parameter entered must be present in the event
 N - match parameters need not be present in the event

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log configuration filter

```
show log configuration filter {<filter name>}
```

Description

Displays the log configuration for the specified filter.

Syntax Description

filter name	Specifies the filter to display.

Default

If no options are specified, the command displays the configuration for all filters.

Usage Guidelines

This command displays the configuration for filters.

Example

The following command displays the configuration for the filter, *myFilter*:

```
show log configuration filter myFilter
```

The following is sample output from this command:

Log Filter Name: myFilter

I/ E	Comp.	Sub-comp.	Condition	Severity
-	-----	-----	-----	-----
I	STP			-----
I	aaa			-----

Include/Exclude: I - Include, E - Exclude

Component Unreg: * - Component/Subcomponent is not currently registered

Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
* - Pre-assigned severities in effect for specified component

```

Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
+ - Debug Severities, but log debug-mode not enabled

```

If Match parameters present:

Parameter Flags: S - Source, D - Destination, (as applicable)
I - Ingress, E - Egress, B - BGP

```
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
                  MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
                  VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
                  L4 - Layer-4 Port #, Num - Number, Str - String
                  Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
                  Proc - Process Name
```

Strict Match : Y - every match parameter entered must be present in the event
N - match parameters need not be present in the event

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log configuration target

```
show log configuration target {console | memory-buffer | nvram | primary-  
msm | primary-node | backup-msm | backup-node | session | syslog  
{<ipaddress> | <ipPort> | vr <vr_name>} {[local0 ... local7]}}
```

Description

Displays the log configuration for the specified target.

Syntax Description

console	Show the log configuration for the console display.
memory-buffer	Show the log configuration for volatile memory.
nvram	Show the log configuration for NVRAM.
primary-msm	Specifies the primary MSM. NOTE: This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM. NOTE: This parameter is available only on modular switches
backup-node	Specifies the backup-node in a stack.
session	Show the log configuration for the current session (including console display).
syslog	Show the configuration for the specified syslog target.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.

Default

If no options are specified, the command displays the configuration for the current session and console display.

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command displays the log configuration for the specified target. The associated filter, severity, match expression, and format is displayed.

Example

The following command displays the log configuration:

```
show log configuration target
```

The following is sample output from this command:

```
Log Target      : memory-buffer
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Debug-Data (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
ion>
  Buffer size  : 1000 messages

Log Target      : nvram
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
ion>

Log Target      : console
  Enabled ?    : no
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Info (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
ion>

Log Target      : primary-msm
  Enabled      : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)

Log Target      : backup-msm
  Enabled      : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

The `local0` ... `local7` keywords were made optional in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show log counters

```
show log counters {<event condition> | [all | <event component>]} {include
| notified | occurred} {severity <severity> {only}}
```

Description

Displays the incident counters for events.

Syntax Description

event condition	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
event component	Specifies that all the events associated with a particular component or subcomponent should be displayed.
include	Specifies if one or more targets should be included in this event.
notified	Specifies the number of times this event has occurred.
occurred	Specifies the number of times this event has occurred since the last clear or reboot.
severity	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed

Default

If `severity` is not specified, then events of all severity are displayed.

Usage Guidelines

This command displays the incident counters for each event specified. Two incident counters are displayed. One counter displays the number of times an event has occurred, and the other displays the number of times that notification for the event was made to the system (an incident record was injected into the system for further processing). Both incident counters reflect totals accumulated since reboot or since the counters were cleared using the `clear log counters` or `clear counters` command, regardless of whether it was filtered or not.

The keywords `include`, `notified`, and `occurred` only display events with non-zero counter values for the corresponding counter.

This command also displays a reference count (the column titled `REF` in the output). The reference count is the number of enabled targets receiving notifications of this event.

See the command `show log` on page 763 for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events
```

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event counters for event conditions of severity debug-summary or greater in the component *STP.InBPDU*:

```
show log counters stp.inbpdu severity debug-summary
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Occurred	In	Notified
-----	-----	-----	-----	-----	-----	-----
STP	InBPDU	Drop	Error	0	Y	0
STP	InBPDU	Ign	Debug-Summary	0	N	0
STP	InBPDU	Mismatch	Warning	0	Y	0

Occurred : # of times this event has occurred since last clear or reboot

Flags : (*) Not all applications responded in time with there count values

In(cluded): Set to Y(es) if one or more targets filter includes this event

Notified : # of times this event has occurred when 'Included' was Y(es)

The following command displays the event counters for the event condition *PDUDrop* in the component *STP.InBPDU*:

```
show log counters "STP.InBPDU.Drop"
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Occurred	In	Notified
-----	-----	-----	-----	-----	-----	-----
STP	InBPDU	Drop	Error	0	Y	0

Occurred : # of times this event has occurred since last clear or reboot

Flags : (*) Not all applications responded in time with there count values

In(cluded): Set to Y(es) if one or more targets filter includes this event

Notified : # of times this event has occurred when 'Included' was Y(es)

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log events

```
show log events [<event condition> | [all | <event component>] {severity
<severity> {only}}] {details}
```

Description

Displays information about the individual events (conditions) that can be logged.

Syntax Description

event condition	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
event component	Specifies that all the events associated with a particular component should be displayed.
severity	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.
details	Specifies that detailed information, including the message format and parameter types, be displayed.

Default

If severity is not specified, then events of all severity are displayed. If detail is not specified, then summary only information is displayed.

Usage Guidelines

This command displays the mnemonic, message format, severity, and parameter types defined for each condition in the event set specified.

See the command [show log on page 763](#) for more information about severity levels.

When the `detail` option is specified, the message format is displayed for the event conditions specified. The message format parameters are replaced by the value of the parameters when the message is generated.

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event conditions of severity debug-summary or greater in the component *STP.InBPDU*:

```
show log events stp.inbpdu severity debug-summary
```


The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Drop	Error	2 total
STP	InBPDU	Ign	Debug-Summary	2 total
STP	InBPDU	Mismatch	Warning	2 total

The following command displays the details of the event condition *PDUTrace* in the component *STP.InBPDU*:

```
show log events stp.inbpdu.pdutrace details
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Trace	Debug-Verbose	2 total 0 - string 1 - string (printf)
		Port=%0%: %1%		

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ports rxerrors

```
show ports {<port_list> | stack-ports <stacking-port-list>} rxerrors {no-
refresh}
```

Description

Displays real-time receive error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
stacking-port-list	Specifies one or more stacking ports or slots. Applies to SummitStack and the Summit family of switches only.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the receive errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, receive error statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time receive error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the `no-refresh` parameter each time you want a snapshot of the port receive errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Receive Error Information. The switch collects the following port receive error information:

- Port Number
- Link Status—The current status of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.

- Receive Bad CRC Frames (RX CRC)—The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- Receive Oversize Frames (RX Over)—The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes.
- Receive Undersize Frames (RX Under)—The total number of frames received by the port that were less than 64 bytes long.
- Receive Fragmented Frames (RX Frag)—The total number of frames received by the port were of incorrect length and contained a bad FCS value.
- Receive Jabber Frames (RX Jabber)—The total number of frames received by the port that was of greater than the support maximum length and had a Cyclic Redundancy Check (CRC) error.
- Receive Alignment Errors (RX Align)—The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- Receive Frames Lost (RX Lost)—The total number of frames received by the port that were lost because of buffer overflow in the switch.

Port Monitoring Display Keys. For information about the available port monitoring display keys, see the `show ports statistics` command.

Example

The following command displays receive error statistics for slot 5, ports 4 through 7 on a modular switch with auto-refresh disabled:

```
show ports 5:4-5:7 rxerrors no-refresh
```

The following is sample output from this command:

```
Port Rx Error monitor
  Port      Link   Rx   Rx   Rx   Rx   Rx   Rx   Rx
         State  Crc  Over Under Frag Jabber Align Lost
=====
   5:4      R      0    0    0    0    0    0    0
   5:5      R      0    0    0    0    0    0    0
   5:6      R      0    0    0    0    0    0    0
   5:7      R      0    0    0    0    0    0    0
=====
Link Status: A-Active, R-Ready, NP-Port not present, L-Loopback
```

The following command displays receive error statistics for all ports on the Summit family of switches with auto-refresh enabled (the default behavior):

```
show ports rxerrors
```

The following is sample truncated output from this command:

```
Port Rx Error Monitor                                     Tue Jul  5 15:07:13 UTC 2005

  Port      Link   Rx   Rx   Rx   Rx   Rx   Rx   Rx
         State  Crc  Over Under Frag Jabber Align Lost
=====
   1         R      0    0    0    0    0    0    0
   2         R      0    0    0    0    0    0    0
   3         R      0    0    0    0    0    0    0
```

4	R	0	0	0	0	0	0	0
5	R	0	0	0	0	0	0	0
6	R	0	0	0	0	0	0	0
7	R	0	0	0	0	0	0	0
8	R	0	0	0	0	0	0	0
9	R	0	0	0	0	0	0	0
10	R	0	0	0	0	0	0	0
11	R	0	0	0	0	0	0	0
12	R	0	0	0	0	0	0	0
13	R	0	0	0	0	0	0	0
14	R	0	0	0	0	0	0	0
15	R	0	0	0	0	0	0	0
16	R	0	0	0	0	0	0	0
17	R	0	0	0	0	0	0	0

```
=====
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

The link status information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the `no-refresh` parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The `no-refresh` parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

The stack-ports option is available only on SummitStack and the Summit family of switches.

show ports statistics

```
show ports {stack-ports <stacking-port-list> | <port_list>} statistics {no-  
refresh}
```

Description

Displays real-time port statistic information. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

stacking-port-list	Specifies one or more stacking slots and ports. Applies to SummitStack and the Summit family of switches only.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the port statistics at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time port statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the `no-refresh` parameter each time you want a snapshot of the port statistics.

Jumbo frame statistics are displayed for switches only that are configured for jumbo frame support.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Statistics. The switch collects the following port statistic information:

- Port Number
- Link Status—The current status of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.

- Transmitted Packet Count (Tx Pkt Count)—The number of packets that have been successfully transmitted by the port.
- Transmitted Byte Count (Tx Byte Count)—The total number of data bytes successfully transmitted by the port.
- Received Packet Count (RX Pkt Count)—The total number of good packets that have been received by the port.
- Received Byte Count (RX Byte Count)—The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- Received Broadcast (RX Bcast)—The total number of frames received by the port that are addressed to a broadcast address.
- Received Multicast (RX Mcast)—The total number of frames received by the port that are addressed to a multicast address.

Port Monitoring Display Keys. Table 18 describes the keys used to control the display that appears if auto-refresh is enabled (the default behavior).

Table 18: Port monitoring display keys with auto-refresh enabled

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc]	Exits from the screen.
O	Clears all counters.

Table 19 describes the keys used to control the display that appears if you auto-refresh is disabled.

Table 19: Port monitoring displays keys with auto-refresh disabled

Key	Description
Q	Exits from the screen.
[Space]	Displays the next page of ports.

Example

The following command displays port statistics for slot 1, ports 1 through 2 on a modular switch with auto-refresh disabled:

```
show ports 1:1-1:2 statistics no-refresh
```

The following is sample output from this command:

```
Port Statistics
Port      Link      Tx Pkt    Tx Byte    Rx Pkt    Rx Byte    Rx      Rx
          Status   Count     Count      Count     Count      Bcast   Mcast
=====
1:1       A         7241      2722608    14482     3968068    0       0
1:2       R          0          0          0          0          0       0
=====
Link Status: A-Active, R-Ready, NP-Port not present, L-Loopback
```

The following command displays port statistics for all ports on the Summit family of switches with auto-refresh enabled (the default behavior):

```
show ports statistics
```

The following is sample truncated output from this command:

```
Port Statistics                                     Tue Jul  5 14:18:34 UTC 2005
Port      Link      Tx Pkt      Tx Byte      Rx Pkt      Rx Byte      Rx      Rx
          Status    Count      Count      Count      Count      Bcast   Mcast
=====
1          R          0          0          0          0          0       0
2          R          0          0          0          0          0       0
3          R          0          0          0          0          0       0
4          R          0          0          0          0          0       0
5          R          0          0          0          0          0       0
6          R          0          0          0          0          0       0
7          R          0          0          0          0          0       0
8          R          0          0          0          0          0       0
9          R          0          0          0          0          0       0
10         R          0          0          0          0          0       0
11         R          0          0          0          0          0       0
12         R          0          0          0          0          0       0
13         R          0          0          0          0          0       0
14         R          0          0          0          0          0       0
15         R          0          0          0          0          0       0
16         R          0          0          0          0          0       0
17         R          0          0          0          0          0       0
=====
Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down ESC->exitPort Statistics
```

History

This command was first available in ExtremeXOS 10.1.

The link status information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the `no-refresh` parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The `no-refresh` parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

The `stack-ports` option is available on SummitStack and the Summit Family of switches only.

show ports txerrors

```
show ports {stack-ports <stacking-port-list> | <port_list>} txerrors {no-
refresh}
```

Description

Displays real-time transmit error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

stacking-port-list	Specifies one or more stacking slot ports for display. Applies to SummitStack and Summit family of switches only.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the transmit errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, error statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time transmit error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the `no-refresh` parameter each time you want a snapshot of the port transmit errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Transmit Error Information. The switch collects the following port transmit error information:

- Port Number
- Link Status—The current status of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.

- Transmit Collisions (TX Coll)—The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- Transmit Late Collisions (TX Late Coll)—The total number of collisions that have occurred after the port's transmit window has expired.
- Transmit Deferred Frames (TX Deferred)—The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- Transmit Errored Frames (TX Errors)—The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).
- Transmit Lost Frames (TX Lost)—The total number of transmit frames that do not get completely transmitted because of buffer problems (FIFO underflow).
- Transmit Parity Frames (TX Parity)—The bit summation has a parity mismatch.

Port Monitoring Display Keys. For information about the available port monitoring display keys, see the `show ports statistics` command.

Example

The following command displays transmit error statistics for slot 5, ports 4 through 7 on a modular switch with auto-refresh disabled:

```
show ports 5:4-5:7 txerrors no-refresh
```

The following is sample output from this command:

```
Port Transmission errors
Port          Link   Tx    Tx    Tx    Tx    Tx    Tx
              State Coll  Late coll  Deferred  Errors  Lost  Parity
=====
   5:4         R      0      0      0      0      0      0
   5:5         R      0      0      0      0      0      0
   5:6         R      0      0      0      0      0      0
   5:7         R      0      0      0      0      0      0
=====
Link Status: A-Active, R-Ready, NP-Port not present, L-Loopback
```

The following command displays transmit error statistics for all ports on a Summit X450 series switch:

```
show ports txerrors
```

The following is sample truncated output from this command:

```
Port Tx Error Monitor                               Tue Jul  5 15:07:13 UTC 2005
Port          Link   Tx    Tx    Tx    Tx    Tx    Tx
              State Coll  Late coll  Deferred  Errors  Lost  Parity
=====
1             R      0      0      0      0      0      0
2             R      0      0      0      0      0      0
3             R      0      0      0      0      0      0
4             R      0      0      0      0      0      0
5             R      0      0      0      0      0      0
6             R      0      0      0      0      0      0
7             R      0      0      0      0      0      0
8             R      0      0      0      0      0      0
9             R      0      0      0      0      0      0
```

10	R	0	0	0	0	0	0
11	R	0	0	0	0	0	0
12	R	0	0	0	0	0	0
13	R	0	0	0	0	0	0
14	R	0	0	0	0	0	0
15	R	0	0	0	0	0	0
16	R	0	0	0	0	0	0
17	R	0	0	0	0	0	0

=====

Link Status : A-Active, R-Ready, NP-Port Not Present, L-Loopback

0->Clear Counters U->page up D->page down ESC->exitPort Tx Error

History

This command was first available in ExtremeXOS 10.1.

The link status information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the `no-refresh` parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The `no-refresh` parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

show ports vlan statistics

```
show ports {<port_list>} vlan statistics {no-refresh}
```

Description

Displays VLAN statistics.



NOTE

This command is available only on the BlackDiamond 12800 series of switches.

Syntax Description

port_list	Specifies one or more ports or slots and ports. Can be one or more port numbers. May be in the form: 1, 2, 3-5, 1:*, 1:5, 1:6-1:8.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.

Default

N/A

Usage Guidelines

This command is used in conjunction with the `configure ports [<port_list>|all] monitor vlan <vlan_name>` command for monitoring port-based VLAN statistics in real time.

Example

The following command displays VLAN information for the VLAN named *corp*:

```
show ports vlan corp vlan stats
```

Displays the vlan statistics in a real time countinuous refresh mode or no-refresh mode.

```
* (debug) BD-12804.31 # show ports 5:1-2 vlan stat
```

```
Port VLAN Statistics                               Wed Mar 28 10:52:59 2007
Port      Vlan      Rx Frames      Rx Byte      Tx Frame      Tx Byte
              Count          Count          Count          Count
=====
5:1      Default  318750522      20400046784      318750588      20400051672
5:2      Default  292811491      18739948736      292811975      18739980504
```

```
0->Clear Counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

show rmon memory

```
show rmon memory {detail | <memoryType>}
```

Description

Displays RMON specific memory usage and statistics.

Syntax Description

detail	Displays detailed information.
memoryType	Specifies the type of memory usage and statistics to display.

Default

N/A.

Usage Guidelines

If you do not specify the `detailed` keyword or enter a specific RMON memory type, the output contains usage information for all memory types.

Example

The following command displays RMON memory statistics:

```
show rmon memory
```

The following is sample output from this command:

```
RMON Memory Information
-----

Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN

Memory Utilization Statistics
-----
      Size      16      32      48      64      80      96      112      128      144      176
208    256    384      5
12    768   1024   2048   4096   8192  16384  18432  40960  64000
-----
-----
Used Blocks   1558      3   2490      1      0      0      0      1      1      0
63444      1   1869
0    311      0      0      0      0      0      0      0      0
rmonEstat      0      0      0      0      0      0      0      0      0
0      0    311
0      0      0      0      0      0      0      0      0
```

```

    rmonOwner 1555      0      0      0      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonHisc    0      0      0      0      0      0      0      0      0      0
0      0 1244
    0      0      0      0      0      0      0      0      0      0
    rmonHist    0      0      0      0      0      0      0      0      0      0
63444      0      0
    0      0      0      0      0      0      0      0      0      0
    rmonAlarm   0      0      0      0      0      0      0      0      0      0
0      0      3
    0      0      0      0      0      0      0      0      0
rmonLogDescription      0      0      0      0      0      0      0      0      0      1
0      0      0      0
    0      0      0      0      0      0      0      0      0
    rmonLog      0      1      0      0      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0      0
    rmonEvent    0      0      0      0      0      0      0      0      1      0      0
0      0      0
    0      0      0      0      0      0      0      0      0
rmonEventDescription      0      1      0      0      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0      0
rmonEventCommunity      0      1      0      0      0      0      0      0      0      0
0      0      0      0
    0      0      0      0      0      0      0      0      0
    rmonCommunity 1      0      0      0      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0      0
    rmonDs      0      0      0      0      0      0      0      0      0      0
0      0      0
    0 311      0      0      0      0      0      0      0
    rmonDbx      0      0 2490      0      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0
    rmonOid      0      0      0      0      0      0      0      0      0      0
0      0 311
    0      0      0      0      0      0      0      0
rmonMdbIndexOid      2      0      0      1      0      0      0      0      0      0
0      0      0
    0      0      0      0      0      0      0      0
    rmonMdbString 0      0      0      0      0      0      0      0      0      0
0      1      0
    0      0      0      0      0      0      0      0

```

The following command displays RMON event statistics:

```
show rmon memory rmonEvent
```

The following is sample output from this command:

RMON Memory Information

Bytes Allocated: 14298032 AllocFailed: 0

Current Memory Utilization Level: GREEN

Memory Utilization Statistics

Memory Statistics for rmonEvent

	Size	16	32	48	64	80	96	112	128	144	176
208	256	384	512	768	1024	2048	4096	8192	16384	18432	40960
64000											

	Allocated	0	0	0	0	0	0	0	1	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0
	AllocatedPeak	0	0	0	0	0	0	0	1	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0
	AllocSuccess	0	0	0	0	0	0	0	1	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0
	FreeSuccess	0	0	0	0	0	0	0	0	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0
	AllocFail	0	0	0	0	0	0	0	0	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0
	FreeFail	0	0	0	0	0	0	0	0	0	0
0	0	0									
0	0	0	0	0	0	0	0	0	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show sflow configuration

```
show sflow {configuration}
```

Description

Displays the current sFlow configuration.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

This command displays the sFlow configuration of your system.

The following fields are displayed:

- Global Status—sFlow is globally enabled or disabled
- Polling interval—How often the hardware is polled for statistics, in seconds
- Sampling rate—Packets are sampled, on average, once for every rate-number of packets
- Maximum cpu sample limit—Maximum number of packets per second sampled before sample throttling takes effect
- Agent IP—IP address inserted into the sFlow data packets to identify the sFlow switch
- Collectors—To which IP address and port, and from which virtual router, the sFlow packets are sent
- Port Status—Enabled or disabled for statistics gathering
- Port Sample-rate—Shows the sampling rate configured for the port and the actual rate if CPU throttling has taken effect
- Port Subsampling factor—See the command `configure sflow ports sample-rate` for details

Example

To display the sFlow configuration on your system, use the following command:

```
show sflow
```

The output from this command is similar to the following:

```
SFLOW Global Configuration
Global Status: enabled
Polling interval: 20
Sampling rate: 8192
Maximum cpu sample limit: 2000
SFLOW Configured Agent IP: 10.203.2.38 Operational Agent IP: 10.203.2.38
Collectors
```


Collector IP 10.201.6.250, Port 6343, VR "VR-Mgmt"

SFLOW Port Configuration

Port	Status	Sample-rate		Subsampling
		Config	/ Actual	factor
1:41	enabled	8192	/ 8192	1
2:40	enabled	1024	/ 1024	1
2:58	enabled	8192	/ 8192	8
2:59	enabled	8192	/ 8192	8

History

This command was first available in an ExtremeXOS 11.0.

Platform Availability

This command is available on all available platforms with an Advanced Edge, a Core, or an Advanced Core license.

show sflow statistics

```
show sflow statistics
```

Description

Displays sFlow statistics.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

This command displays sFlow statistics for your system.

The following fields are displayed:

- Received frames—Number of frames received on sFlow enabled ports
- Sampled Frames—Number of packets that have been sampled by sFlow
- Transmitted Frames—Number of UDP packets sent to remote collector(s)
- Broadcast Frames—Number of broadcast frames received on sFlow enabled ports
- Multicast Frames—Number of multicast frames received on sFlow enabled ports
- Packet Drops—Number of samples dropped

Example

To display sFlow statistics for your system, use the following command:

```
show sflow statistics
```

The output from this command is similar to the following:

```
SFLOW Statistics
```

```
Received frames      : 1159044921
Sampled Frames       : 104944
Transmitted Frames   : 10518
Broadcast Frames     : 0
Multicast Frames     : 1055652
Packet Drops         : 0
```

History

This command was first available in an ExtremeXOS 11.0.

Platform Availability

This command is available on all available platforms with an Advanced Edge, a Core, or an Advanced Core license.

show temperature

```
show temperature
```

Description

Depending on the platform, this command displays the current temperature of the I/O modules, management modules, power supply controllers, XGM-2xn card, and the switch.

On a stack, the command displays the current temperature of the modules in each slot.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different temperature information might be displayed.

Modular Switches Only. Use this command to display the temperature in Celsius and the current status of the following installed components in the switch:

- Management modules (MSM)
- I/O modules
- Power controllers

The switch monitors the temperature of each component and generates a warning if the temperature exceeds the normal operating range. If the temperature exceeds the minimum/maximum limits, the switch shuts down the overheated module.

Summit Family of Switches Only. Use this command to display the temperature in Celsius and the current status of the following components:

- Switch
- XGM-2xn card

SummitStack Only. Use this command to display the temperature in Celsius and the current status of the following components:

- All switches in the stack
- XGM-2xn cards if present

The switch monitors its temperature and generates a warning if the temperature exceeds the normal operating range. If the temperature exceeds the maximum limit, the `show switch` output indicates the

switch in an OPERATIONAL (Overheat) mode, and the `show temperature` output indicates an error state due to overheat.

Displaying the Temperature of Other Installed Components—Modular Switches Only. You can also view the temperature of the power supplies and the fan trays in the switch.

To view the temperature of the power supplies installed in a modular switch, use the following command:

```
show power {<ps_num>} {detail}
```

To view the temperature of the fan trays installed in the BlackDiamond 10808 switch, use the following command:

```
show fans {detail}
```

Example

Depending on the platform, the following command displays the temperature of various switch components:

```
show temperature
```

The following is sample output from a modular switch:

Field	Replaceable Units	Temp (C)	Status
Slot-1	: 10G6X	36.37	Normal
Slot-2	: G60X	35.31	Normal
Slot-3	:		
Slot-4	:		
Slot-5	:		
Slot-6	: G60X	34.68	Normal
Slot-7	: G60X	34.31	Normal
Slot-8	:		
MSM-A	: MSM-1XL	31.37	Normal
MSM-B	: MSM-1XL	29.75	Normal
PSUCTRL-1	:		
PSUCTRL-2	:	29.00	Normal

Temp Range: -10.00 (Min), 0.00-50.00 (Normal), 60.00 (Max)

The following is sample output from a Summit X450 series switch:

Field	Replaceable Units	Temp (C)	Status
Switch	: SummitX450-24x	32.00	Normal
XGM-2xn-1	:		

Temp Range: -10.00 (Min), 0.00-60.50 (Normal), 62.90 (Max)

The following is sample output from a SummitStack of 8 nodes:

```
* Slot-4 Stack.32 # show temperature
Field Replaceable Units          Temp (C)   Status
-----
Slot-1      : SummitX450-24x      40.00    Normal
Slot-2      : SummitX450-24t      38.00    Normal
Slot-3      : X450e-48p           38.00    Normal
Slot-4      : X450e-24p           42.00    Normal
Slot-5      : X450a-24t           43.00    Normal
Slot-6      : X450a-48t           40.50    Normal
Slot-7      : X450a-24x           46.00    Normal
Slot-8      : X450a-24t           35.50    Normal
XGM2-2xn-1  :
10GBase-1   :
10GBase-2   :
```

Temp Range: -10.00 (Min), 0.00-60.50 (Normal), 62.90 (Max)

History

This command was first available in an ExtremeXOS 10.1.

Information about the power controller(s), a component status column, and the minimum, normal, and maximum temperature ranges of the components was added to the `show temperature` output in ExtremeXOS 11.0.

Information about the XGM-2xn card was added to the `show temperature` output in ExtremeXOS 11.2.

Support for stacking was added to `show temperature` output in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show version

```
show version {detail | process <name> | images {partition <partition>} {msm
<slotid>} }
```

Description

Displays the hardware serial and version numbers, the software version currently running on the switch, and (if applicable) the software version running on the modules and power controllers.

Syntax Description

detail	Specifies display of slot board name and chassis or platform name.
process	Specifies display of all of the processes on the switch.
name	Specifies display of a specific process on the switch.
images	Specifies the display of installed images.
partition	Specifies display of a specific partition (primary or secondary).
slotid	Specifies display of an MSM in a specific slot (A or B).
	NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

The following describes the information displayed when you execute the `show version` or `show version detail` commands:

- **Part Number**—A collection of numbers and letters that make up the part number of the switch and when applicable the hardware components installed in a modular switch.
- **Serial Number**—A collection of numbers and letters that make up the serial number of the switch and when applicable the hardware components installed in a modular switch.



NOTE

For information about the physical location of the serial number on your switch, refer to the section that describes your specific switch model in the Extreme Networks Consolidated XOS Hardware Guide.

- **Image**—The ExtremeXOS software version currently running on the switch. If you have two software images downloaded on the switch, only the currently running ExtremeXOS version information is displayed. The information displayed includes the major version number, minor version number, a specific patch release, and the build number. The software build date is also displayed.
- **BootROM**—The BootROM version currently running on the switch.

- uC—The micro controller firmware version currently running on the module (BlackDiamond 12804 switch only).
- FPGA—The field-programmable gate array firmware version currently running on the module (BlackDiamond 12804 switch only).

Depending on the model of your switch and the software running on your switch, different version information may be displayed.

**NOTE**

The information displayed does not include the I/O version number or the BlackDiamond 8800 series switch. The I/O version number includes the major, minor, and I/O version number, not the patch and build numbers.

If you use the `process` option, you will see the following information about the processes running on the switch:

- Card—The location (MSM) where the process is running on a modular switch.
- Process Name—The name of the process.
- Version—The version number of the process.
- BuiltBy—The name of the software build manager.
- Link Date—The date the executable was linked.

Example

The following command displays the hardware and software versions currently running on the switch:

```
show version
```

The following is sample output from a BlackDiamond 10808 switch:

```
Chassis      : 804300-00-09 0438F-00850 Rev 9.0
Slot-1       : 804403-00-08 0405F-00063 Rev 8.0 BootROM: 1.3.0.0      IMG: 11.4.0.18
Slot-2       : 804406-00-02 0511F-00605 Rev 2.0
Slot-3       : 804406-00-02 0511F-00602 Rev 2.0
Slot-4       :
Slot-5       :
Slot-6       :
Slot-7       :
Slot-8       :
MSM-A        : 804301-00-09 0434F-00125 Rev 9.0 BootROM: 1.0.1.5      IMG: 11.4.0.18
MSM-B        : 804301-00-09 0437F-00026 Rev 9.0
PSUCTRL-1    : 704025-00-06 0438F-00089 Rev 6.0 BootROM: 5.3
PSUCTRL-2    : 704025-00-06 0438F-00096 Rev 6.0 BootROM: 5.3

Image        : ExtremeXOS version 11.4.0.18 v1140b18 by release-manager
               on Thu Jan 19 18:25:51 PST 2006
BootROM      : 1.0.1.5
```


The following is sample output from the BlackDiamond 8810 switch (the output from the BlackDiamond 8806 is similar):

```
Chassis      : 800129-00-02 04344-00039 Rev 2.0
Slot-1      : 800114-00-04 04364-00021 Rev 4.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
Slot-2      : 800115-00-02 04344-00006 Rev 2.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
Slot-3      : 800113-00-04 04354-00031 Rev 4.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
Slot-4      :
Slot-5      : 800112-00-03 04334-00040 Rev 3.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
Slot-6      : 800112-00-03 04334-00004 Rev 3.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
Slot-7      :
Slot-8      :
Slot-9      :
Slot-10     :
MSM-A       : 800112-00-03 04334-00040 Rev 3.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
MSM-B       : 800112-00-03 04334-00004 Rev 3.0 BootROM: 1.0.1.7    IMG: 11.4.0.18
PSUCTRL-1   : 450117-00-01 04334-00021 Rev 1.0 BootROM: 2.13
PSUCTRL-2   : 450117-00-01 04334-00068 Rev 1.0 BootROM: 2.13

Image       : ExtremeXOS version 11.4.0.18 v1140b18 by release-manager
              on Thu Jan 19 19:01:18 PST 2006
BootROM     : 1.0.1.7
```

The following is sample output from the BlackDiamond 12804 switch:

```
Chassis      : 804023-00-01 0515E-00999 Rev 1.0
Slot-1      : 804019-00-01 05204-00004 Rev 1.0 uC: 2.00 FPGA: 5.07
Slot-2      : 804019-00-02 05404-00055 Rev 2.0 uC: 2.00 FPGA: 5.07
Slot-5      : 804022-00-02 05334-00005 Rev 2.0 uC: 2.00 FPGA: 5.07
Slot-6      : 804019-00-01 05234-00026 Rev 1.0 uC: 2.00 FPGA: 5.07
MSM-A       : 804017-00-02 05394-00007 Rev 2.0 BootROM: 0.3.0.9
              IMG: 11.4.0.18 uC: 1.07 FPGA: 3.1f
MSM-B       : 804017-00-02 05394-00011 Rev 2.0 BootROM: 0.3.0.9
              IMG: 11.4.0.18 uC: 2.00
PSUCTRL-1   : 700087-00-03 04454-00007 Rev 3.0 BootROM: 2.13
PSUCTRL-2   :

Image       : ExtremeXOS version 11.4.0.18 v1140b18 by release-manager
              on Thu Jan 19 18:17:02 PST 2006
BootROM     : 0.3.0.9
```

The following is sample output from the Summit X450 series switch:

```
Switch      : 800132-00-01 0452G-01524 Rev 1.0 BootROM: 1.0.0.4    IMG:
XGM-2xn-1   :

Image       : ExtremeXOS version 11.2.0.14 v1120b14 by release-manager
              on Thu Mar 31 09:33:26 PST 2005
BootROM     : 1.0.0.4
```

This is a sample from a SummitStack:

```
Slot-1      : 800152-00-04 0624G-01015 Rev 4.0 BootROM: 1.0.2.0    IMG: 12.0.0.4
Slot-2      : 800192-00-01 0620G-00020 Rev 1.0 BootROM: 1.0.1.9    IMG: 12.0.0.4
Slot-3      : 800163-00-04 0630G-00804 Rev 4.0 BootROM: 1.0.2.0    IMG: 12.0.0.4
```

```

Slot-4      : 800153-00-04 0630G-00672 Rev 4.0 BootROM: 1.0.2.0    IMG: 12.0.0.4
Slot-5      :
Slot-6      :
Slot-7      :
Slot-8      :
XGM2-1      :

```

```

Image       : ExtremeXOS version 12.0.0.4 branch-fixes_v1200b4 by mmroz
              on Thu Mar 15 10:55:58 EDT 2007
BootROM     : 1.0.2.0

```

Using the `process` option of the `show version` command produces output similar to the following on a modular switch:

Card Process Name	Version	BuiltBy	Link Date
MSM-A aaa	3.0.0.2	release-manager	Thu Mar 31 09:23:54 PST 2005
MSM-A acl	3.0.0.2	release-manager	Thu Mar 31 09:26:46 PST 2005
MSM-A bgp	3.0.0.2	release-manager	Thu Mar 31 09:27:54 PST 2005
MSM-A cfgmgr	3.0.0.21	release-manager	Thu Mar 31 09:23:42 PST 2005
MSM-A cli	3.0.0.22	release-manager	Thu Mar 31 09:23:34 PST 2005
MSM-A devmgr	3.0.0.2	release-manager	Thu Mar 31 09:23:22 PST 2005
MSM-A dirser	3.0.0.2	release-manager	Thu Mar 31 09:24:02 PST 2005
MSM-A eaps	3.0.0.8	release-manager	Thu Mar 31 09:26:34 PST 2005
MSM-A edp	3.0.0.2	release-manager	Thu Mar 31 09:25:56 PST 2005
MSM-A elrp	3.0.0.1	release-manager	Thu Mar 31 09:25:14 PST 2005
MSM-A ems	3.0.0.2	release-manager	Thu Mar 31 09:35:08 PST 2005
MSM-A epm	3.0.0.3	release-manager	Thu Mar 31 09:23:11 PST 2005
MSM-A esrp	3.0.0.4	release-manager	Thu Mar 31 09:26:23 PST 2005
....			

The following is sample output from the Summit X450 series switch:

Process Name	Version	BuiltBy	Link Date
aaa	3.0.0.2	release-manager	Thu Mar 31 09:34:17 PST 2005
acl	3.0.0.2	release-manager	Thu Mar 31 09:38:44 PST 2005
bgp	Not Started	Unknown	Unknown
cfgmgr	3.0.0.21	release-manager	Thu Mar 31 09:33:58 PST 2005
cli	3.0.0.22	release-manager	Thu Mar 31 09:33:45 PST 2005
cna	3.1.0.1	release-manager	Thu Mar 31 09:49:28 PST 2005
devmgr	3.0.0.2	release-manager	Thu Mar 31 09:33:26 PST 2005
dirser	3.0.0.2	release-manager	Thu Mar 31 09:34:31 PST 2005
dosprotect	3.0.0.1	release-manager	Thu Mar 31 09:48:58 PST 2005
eaps	3.0.0.8	release-manager	Thu Mar 31 09:38:25 PST 2005
edp	3.0.0.2	release-manager	Thu Mar 31 09:37:24 PST 2005
elrp	3.0.0.1	release-manager	Thu Mar 31 09:36:27 PST 2005
ems	3.0.0.2	release-manager	Thu Mar 31 09:50:40 PST 2005
epm	3.0.0.3	release-manager	Thu Mar 31 09:33:08 PST 2005
esrp	3.0.0.4	release-manager	Thu Mar 31 09:38:07 PST 2005
etmon	1.0.0.1	release-manager	Thu Mar 31 09:47:16 PST 2005
....			

If you specify the `name` option, only the process you select is displayed.

Using the `images` option in the `show version` command produces output similar to the following on a modular switch:

Card	Partition	Installation Date	Version	Name
MSM-A	primary	Wed Jun 30 22:30:22 UTC 2004	11.0.0.24	bd10K-11.0.0.24.xos
MSM-A	primary	Thu Jul 1 03:29:41 UTC 2004	11.0.0.24	bd10K-11.0.0.24-ssh.xmod
MSM-A	secondary	Tue Jun 29 06:09:26 UTC 2004	11.0.0.23	bd10K-11.0.0.23.xos
MSM-A	secondary	Tue Jun 29 06:29:14 UTC 2004	11.0.0.23	bd10K-11.0.0.23-ssh.xmod

The following is sample output from the Summit X450 series switch:

Card	Partition	Installation Date	Version	Name	Branch
Switch	primary	Wed Feb 4 06:36:30 UTC 2004	11.2.0.14	summitx450-11.2.0.14.xos	v1120b14
Switch	primary	Wed Feb 4 07:05:35 UTC 2004	11.2.0.14	summitx450-11.2.0.14-ssh.xmod	v1120b14
Switch	primary	Wed Feb 4 07:09:13 UTC 2004	11.2.0.14	summitx450-11.2.0.14-cna.xmod	v1120b14
Switch	secondary	Mon Jan 19 01:52:28 UTC 2004	11.2.0.10	summitx450-11.2.0.10.xos	v1120b10
Switch	secondary	Wed Jan 21 06:50:42 UTC 2004	11.2.0.10	summitx450-11.2.0.10-ssh.xmod	v1120b10
Switch	secondary	Wed Jan 21 06:51:13 UTC 2004	11.2.0.10	summitx450-11.2.0.10-cna.xmod	v1120b10

If you specify the `partition` option, only images on the specified partition is shown.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vlan statistics

```
show vlan <vlan_name> statistics {no-refresh}
```

Description

Displays VLAN statistics.



NOTE

This command is available only on the BlackDiamond 12800 series of switches.

Syntax Description

vlan_name	Specifies to display VLAN statistics from the VLAN with this name.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.

Default

N/A

Usage Guidelines

The user must configure all the ports on a VLAN to monitor the VLAN statistics using the command `configure ports all monitor vlan <vlan_name>`.

Example

The following command displays VLAN statistics:

```
* (debug) BD-12804.17 # show vlan statistics no-refresh
```

Vlan	Rx Total Frames	Rx Byte Count	Tx Total Frames	Tx Byte Count
Default	30251013	7326296	22034	901840

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

unconfigure log filter

```
unconfigure log filter <filter name>
```

Description

Resets the log filter to its default values; removes all filter items.

Syntax Description

filter name	Specifies the log filter to unconfigure.

Default

N/A.

Usage Guidelines

If the filter name specified is *DefaultFilter*, this command restores the configuration of *DefaultFilter* back to its original settings.

If the filter name specified is not *DefaultFilter*, this command sets the filter to have no events configured and therefore, no incidents will pass. This is the configuration of a newly created filter that was not copied from an existing one.

See the `delete log filter` command for information about deleting a filter.

Example

The following command sets the log filter myFilter to stop passing any events:

```
unconfigure log filter myFilter
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available all platforms.

unconfigure log target format

```
unconfigure log target [console | memory-buffer | nvram | session | syslog
[all | <ipaddress> | <ipPort> {vr <vr_name>} [local0 ... local7]]] format
```

Description

Resets the log target format to its default values.

Syntax Description

console	Specifies the console display format.
memory-buffer	Specifies the switch memory buffer format.
nvram	Specifies the switch NVRAM format.
session	Specifies the current session (including console display) format.
syslog	Specifies a syslog target format.
all	Specifies all remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
local0 ... local7	Specifies the local syslog facility.
format	Specifies that the format for the target will be reset to the default value.

Default

When a target format is unconfigured, it is reset to the default values.

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- severity—on
- event-name—condition
- host-name—off
- sequence-number—off
- process-name—off
- process-slot—on (modular switches only)
- process-id—off
- source-line—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd

- severity—on
- event-name—none
- host-name—off
- sequence-number—off
- process-name—off
- process-slot—on (modular switches only)
- process-id—off
- source-line—off

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Use this command to reset the target format to the default format.

Example

The following command sets the log format for the target `session` (the current session) to the default:

```
unconfigure log target session format
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure ports monitor vlan

```
unconfigure ports [<port_list>|all] monitor vlan <vlan_name>
```

Description

Stops counting VLAN statistics on a port or group of ports.



NOTE

This command is available only on the BlackDiamond 12800 series switches.

Syntax Description

port_list	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
vlan_name	Specifies a VLAN name.

Default

N/A

Usage Guidelines

None

Example

The following command removes monitoring for ports on the VLAN named *accounting*:

```
unconfigure ports 8:1-8:6 monitor vlan accounting
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 12800 series platforms.

unconfigure sflow

```
unconfigure sflow
```

Description

Resets all the sFlow values to the default values.

Syntax Description

This command has no arguments or variables

Default

The default values for sFlow are as follows:

- sFlow agent IP address—0.0.0.0
- sampling frequency—sample one every 8196 packets
- polling interval—20 seconds
- maximum CPU sample limit—2000 samples per second

sFlow is unconfigured and disabled on all ports.

Usage Guidelines

This command resets sFlow values to the default values, and removes any port configurations, and any sFlow collectors configured on the switch.

Example

The following command unconfigures sFlow:

```
unconfigure sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

unconfigure sflow agent

```
unconfigure sflow agent
```

Description

Resets the sFlow agent's IP address to the default value.

Syntax Description

This command has no arguments or variables.

Default

The default IP address is 0.0.0.0.

Usage Guidelines

This command resets the sFlow agent IP address to its default value.

Example

The following command resets the agent IP back to the management IP address:

```
unconfigure sflow agent
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

unconfigure sflow collector

```
unconfigure sflow collector {ipaddress} <ip-address> {port <udp-port-  
number>} {vr <vrname>}
```

Description

Unconfigures the sFlow collector.

Syntax Description

ip-address	Specifies the IP address of the collector to reset.
udp-port-number	Specifies the UDP port.
vrname	Specifies which virtual router.
	NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—VR-Mgmt (previously called VR-0).

Usage Guidelines



NOTE

The BlackDiamond 8800 series, SummitStack, and the Summit family switches do not support user-created VRs.

This command allows you to reset the specified sFlow collector parameters to the default values.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow collector` will reset the collector parameters to the default.

Example

The following command removes the collector at IP address 192.168.57.1:

```
unconfigure sflow collector ipaddress 192.168.57.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

unconfigure sflow ports

```
unconfigure sflow ports <port_list>
```

Description

Removes the specified ports from the sFlow configuration, and stops sampling them.

Syntax Description

port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

This command removes the specified ports from the sFlow configuration, and stops sampling them.

Example

The following command unconfigures sFlow on the ports 2:5-2:7:

```
unconfigure sflow ports 2:5-2:7
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

upload log

```
upload log <ipaddress> {vr <vr_name>} <filename> {messages [memory-buffer |
nvr] {events {<event-condition> | <event_component>}}} {<severity>
{only}} {match <regex>} {chronological}
```

Description

Uploads the current log messages to a TFTP server.

Syntax Description

ipaddress	Specifies the ipaddress of the TFTP server.
vr_name	Specifies the virtual router that can reach the TFTP server. NOTE: The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.
filename	Specifies the file name for the log stored on the TFTP server.
messages	Specifies the location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory.
nvr	Show messages stored in NVRAM
events	Show event messages.
event-condition	Specifies the event condition to display.
event-component	Specifies the event component to display.
severity	Specifies the minimum severity level to display (if the keyword <code>only</code> is omitted).
only	Specifies that only the specified severity level is to be displayed.
regex	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies uploading log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer
- severity—none (displays everything stored in the target)
- match—no restriction
- chronological—if not specified, show messages in order from newest to oldest

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

This command is similar to the `show log` command, but instead of displaying the log contents on the command line, this command saves the log to a file on the TFTP server you specify. For more details on most of the options of this command, see the command [show log on page 763](#).

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a local or remote file, remember the requirements listed above.

Example

The following command uploads messages with a critical severity to the filename *switch4critical.log* on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4critical.log critical
```

The following command uploads messages with warning, error, or critical severity to the filename *switch4warn.log* on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4warn.log warning
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

12 VLAN Commands

This chapter describes commands for:

- Creating and deleting VLANs and performing basic VLAN configuration
- Enabling and disabling VLANs
- Defining protocol filters for use with VLANs

VLANs can be created according to the following criteria:

- Physical port—A port-based VLAN consists of a group of one or more ports on the switch. A port can be a member of only one port-based VLAN, and is by default a member of the VLAN named Default.
- 802.1Q tag—Tagging is most commonly used to create VLANs that span switches.
- Ethernet, LLC SAP, or LLC/SNAP Ethernet protocol type—Protocol-based VLANs are most often used in situations where network segments contain hosts running multiple protocols.
- A combination of these criteria.

Beginning with ExtremeXOS 11.2, the software supports using IPv6 addresses, in addition to IPv4 addresses.

Beginning with ExtremeXOS version 11.4, the software anticipates the emerging IEEE 802.1ah Backbone Bridge standard, also known as MAC-in-MAC tunneling. This standard allows Internet Service Providers (ISPs) to use Ethernet to create a separate backbone over which the subscriber's frames are transported.

clear mac-binding

```
clear mac-binding {bvlan <bvlan-name> {<bridge-mac>} | svlan <svlan-name>
<customer-mac>}
```

Description

Clears the dynamically learned MAC binding database used in IEEE 802.1ah, or MAC-in-MAC tunneling. This database shows the association between the remote bridge MAC addresses and the customer host MAC addresses behind that bridge.

Syntax Description

bvlan-name	Enter the name of the BVLAN you want to clear.
bridge-mac	Enter the BVLAN MAC address you want to clear.
svlan-name	Enter the name of the SVLAN you want to clear.
customer-mac	Enter the customer MAC address you want to clear.

Default

N/A.

Usage Guidelines

Use this command to clear the database that binds the BVLAN MAC address to the SVLAN, or customer, MAC address.

To display the MAC addresses, issue the `show mac-binding` command.

Example

The following command clears the MAC binding table on the switch:

```
clear mac-binding
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan add ports tagged

```
configure bvlan <bvlan-name> add ports <port_list> tagged
```

Description

Adds ports to the specified BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN to which you are adding ports.

Default

N/A.

Usage Guidelines

Only tagged ports can be added to the BVLAN.

You cannot configure Connectivity Fault Management (CFM) on ports added to the BVLAN, or apply ACLs to those ports.

(You can configure an SVLAN and a BVLAN on the same physical port as long as there is no membership relationship between the specific SVLAN and BVLAN. For example, if SVLAN s1 is in BVLAN b1, they *cannot* share the same physical port.)

Example

The following command assigns the ports 1:1 to 1:5 to the backbone VLAN (BVLAN) black and tags those ports:

```
configure bvlan black add ports 5:1-5 tagged
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan add svlan

```
configure bvlan <bvlan-name> add svlan <svlan-name>
```

Description

Adds the specified SVLAN to the specified BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN to which you are adding a SVLAN.
svlan-name	Enter the name of the SVLAN you are adding to the specified BVLAN.

Default

N/A.

Usage Guidelines

Use this command to add SVLANs to the BVLAN.

A BVLAN may have multiple SVLANs. An SVLAN can be in only one BVLAN and must have a unique service ID within that BVLAN.



NOTE

When you add an SVLAN to a BVLAN, they cannot have overlapping ports. You cannot have the same ports on a specific SVLAN that you add to a specific BVLAN.

Example

The following command adds the SVLAN white to the BVLAN black:

```
configure bvlan black add svlan white
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan delete ports

```
configure bvlan <bvlan-name> delete ports [all | <port_list>]
```

Description

Deletes ports from the specified BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN from which you are deleting ports.
all	Deletes all ports from the BVLAN.
port_list	Specify the ports you want to delete from the VVLAN.

Default

N/A.

Usage Guidelines

Use this command to delete ports from the BVLAN.

Example

The following command deletes the ports 1:1 and 1:2 from the backbone VLAN (BVLAN) black:

```
configure bvlan black delete ports 5:1-2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan delete svlan

```
configure bvlan <bvlan-name> delete svlan [<svlan-name> | all]
```

Description

Deletes the specified SVLAN from the specified BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN from which you are deleting a SVLAN.
svlan-name	Enter the name of the SVLAN you are deleting from the specified BVLAN.
all	Deletes all SVLANs from the specified BVLAN.

Default

N/A.

Usage Guidelines

Use this command to delete SVLANs from the BVLAN.

Example

The following command deletes the SVLAN white from the BVLAN black:

```
configure bvlan black delete svlan white
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan ethertype

```
configure bvlan ethertype <value>
```

Description

The default Ethernet value for 802.1ah, or MAC-in-MAC frames is 0x88B5. This command changes the Ethernet value for 802.1ah frames.

Syntax Description

value	Enter the value in hexadecimal of the Ethernet type you want to use for 802.1ah frames.

Default

N/A.

Usage Guidelines

If you are going to change this, Extreme Networks recommends setting this value prior to any other 802.1ah configuration. Changing the 802.1ah may cause the system to flush both the FDB and the MAC-binding tables.

To display the 802.1ah Ethernet type, issue the `show vman` command.

Example

The following command changes the Ethernet type of a MAC-in-MAC, or 802.1ah, frame from the default of 0x88B5 to 0x88A2:

```
configure bvlan etherthype 0x88A2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure bvlan tag

```
configure bvlan <bvlan-name> tag <tag>
```

Description

Assigns the tag value for the specified BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN to which you are assigning a tag.
tag	Assign a value between 2 and 4094.

Default

N/A.

Usage Guidelines

The BVLAN must be tagged. You can have up to 4094 BVLANs on each switch.

Example

The following command assigns the tag 100 to the backbone VLAN (BVLAN) black:

```
configure bvlan black tag 100
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure protocol add

```
configure protocol <name> add [etype | llc | snap] <hex> {[etype | llc | snap] <hex>}
```

Description

Configures a user-defined protocol filter.

Syntax Description

name	Specifies a protocol filter name.
hex	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: • The Ethernet protocol type taken from a list maintained by the IEEE. • The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). • The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- etype – IEEE Ethertype.
- llc – LLC Service Advertising Protocol.
- snap – Ethertype inside an IEEE SNAP packet encapsulation.

A maximum of 15 protocol filters, each containing a maximum of six protocols, can be defined.

The protocol filter must already exist before you can use this command: use the [create bvlan](#) command to create the protocol filter.

No more than seven protocols can be active and configured for use.

Example

The following command configures a protocol named Fred by adding protocol type LLC SAP with a value of FFEF:

```
configure protocol fred add llc 0xfeff
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure protocol delete

```
configure protocol <name> delete [etype | llc | snap] <hex> {[etype | llc | snap] <hex>} ...
```

Description

Deletes the specified protocol type from a protocol filter.

Syntax Description

name	Specifies a protocol filter name.
hex	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: - The Ethernet protocol type taken from a list maintained by the IEEE. - The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). - The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- etype – IEEE Ethertype.
- llc – LLC Service Advertising Protocol.
- snap – Ethertype inside an IEEE SNAP packet encapsulation.

Example

The following command deletes protocol type LLC SAP with a value of FEFF from protocol *fred*:

```
configure protocol fred delete llc feff
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure svlan add ports

```
configure svlan <svlan-name> add ports <port_list> {tagged | untagged}
```

Description

Adds ports to the specified SVLAN; ports may be tagged or untagged.

Syntax Description

svlan-name	Enter the name of the SVLAN you are adding ports to.
port_list	Specify the ports you want to add to the SVLAN.
tagged	Add this parameter if you to add the port to the specified SVLAN as a tagged port.
untagged	Add this parameter to add untagged ports to the specified SVLAN. This is the default value; if you do not specify, the ports are added untagged.

Default

Untagged.

Usage Guidelines

The ports in the SVLAN may be tagged or untagged.

You cannot configure Connectivity Fault Management (CFM) on ports added to the SVLAN, or apply any other ACLs.

(You can configure an SVLAN and a BVLAN on the same physical port as long as there is no membership relationship between the specific SVLAN and BVLAN. For example, if SVLAN s1 is in BVLAN b1, they *cannot* share the same physical port.)

Example

The following command assigns port 1:5 as a tagged port to the SVLAN white:

```
configure svlan white add ports 1:5 tagged
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure svlan delete ports

```
configure svlan <svlan-name> delete ports [all | <port_list>]
```

Description

Deletes ports from the specified SVLAN.

Syntax Description

svlan-name	Enter the name of the SVLAN from which you are deleting ports.
all	Deletes all ports from the SVLAN.
port_list	Specify the ports you want to delete from the SVLAN.

Default

N/A.

Usage Guidelines

Use this command to delete ports from the SVLAN.

Example

The following command deletes port 3:5 from the SVLAN white:

```
configure svlan white delete ports 3:5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure svlan isid

```
configure svlan <svlan-name> isid <isid>
```

Description

This optional command configures the ISID on the 802.1ah, or MAC-in-MAC, frame, which is used to map the S-tag on the original vMAN packet to the I-tag on the 802.1ah frame. Each ISID must be unique in a BVLAN. If you do not configure this value, the system uses the existing value on the S-tag.

Syntax Description

svlan-name	Enter the name of the SVLAN to which you are assigning an ISID.
isid	Enter a value between 0 and 16,777,215 to assign an ISID to the specified SVLAN.

Default

N/A.

Usage Guidelines

This is an optional command when configuring MAC-in-MAC, or 802.1ah, tunnels.

This value is used by the system to map the original S-tag on the vMAN, or 802.1ad, packet to the I-tag on the MAC-in-MAC, or 802.1ah, frame. This value is used in the MAC-binding tables that hold this mapping information.

If you do not configure this value, the system uses the existing value on the S-tag for the map.



NOTE

You must configure an ISID if you have duplicate S-tag values on the same BVLAN ingress port or if you need more than 4094 S-tag values for one BVLAN.

Example

The following command assigns an I-SID value of 250 to the SVLAN white:

```
configure svlan white isid 250
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure svlan tag

```
configure svlan <svlan-name> tag <tag>
```

Description

Assigns a tag to the SVLAN, which must be tagged.

Syntax Description

svlan-name	Enter the name of the SVLAN you are tagging.
tag	Enter a value between 2 and 4094 to tag the specified SVLAN.

Default

N/A.

Usage Guidelines

The SVLAN must be tagged.

You can use a tag previously assigned to another SVLAN, as long as they are not on the same ingress port.

Example

The following command assigns a tag value of 200 to the SVLAN white:

```
configure svlan white tag 200
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

configure vlan add ports

```
configure vlan <vlan_name> add [ports <port_list> | all] {tagged |
untagged} {stpd <stpd_name> {dot1d | emistp | pvst-plus}}
```

Description

Adds one or more ports in a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
port_list	Specifies a list of ports or slots and ports.
all	Specifies all ports.
tagged	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
stpd_name	Specifies an STP domain name.
dot1d emistp pvst-plus	Specifies the BPDU encapsulation mode for these STP ports.

Default

Untagged.

Usage Guidelines

The VLAN must already exist before you can add (or delete) ports: use the [create vlan](#) command to create the VLAN.

If the VLAN uses 802.1Q tagging, you can specify tagged or untagged port(s). If the VLAN is untagged, the ports cannot be tagged.

Untagged ports can only be a member of a single VLAN. By default, they are members of the default VLAN (named *Default*). In order to add untagged ports to a different VLAN, you must first remove them from the default VLAN. You do not need to do this to add them to another VLAN as tagged ports. If you attempt to add an untagged port to a VLAN prior to removing it from the default VLAN, you see the following error message:

```
Error: Protocol conflict when adding untagged port 1:2. Either add this port as
tagged or assign another protocol to this VLAN.
```

As each port can belong to only one virtual router, ports within one VLAN must all be in the same virtual router.



NOTE

All ports and VLANs are in VR-Default in the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches and cannot be moved.

Refer to [Chapter 18, “STP Commands,”](#) for more information on configuring Spanning Tree Domains.

**NOTE**

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Beginning with ExtremeXOS version 11.4, the system returns the following message if the ports you are adding are already EAPS primary or EAPS secondary ports:

```
WARNING: Make sure Vlan1 is protected by EAPS, Adding EAPS ring ports to a VLAN
could cause a loop in the network.
Do you really want to add these ports? (y/n)
```

Example

The following command assigns tagged ports 1:1, 1:2, 1:3, and 1:6 to a VLAN named *accounting*:

```
configure vlan accounting add ports 1:1, 1:2, 1:3, 1:6 tagged
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan delete ports

```
configure vlan <vlan_name> delete ports [all | <port_list>]
```

Description

Deletes one or more ports in a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all ports.
port_list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes ports 1:1, 1:2, 4:3, and 5:6 on a modular switch from a VLAN named *accounting*:

```
configure accounting delete port 1:1, 1:2, 4:3, 5:6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan ipaddress

```
configure vlan <vlan_name> ipaddress [<ipaddress> {<ipNetmask>} |
ipv6-link-local | {eui64} <ipv6_address_mask>]
```

Description

Assigns an IPv4 address and an optional subnet mask or an IPv6 address to the VLAN. Beginning with ExtremeXOS version 11.2, you can specify IPv6 addresses. You can assign either an IPv4 address, and IPv6 address, or both to the VLAN. Beginning with ExtremeXOS software version 11.3, you can use this command to assign an IP address to a specified vMAN and enable multicasting on that vMAN.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IPv4 address.
ipNetmask	Specifies an IPv4 subnet mask in dotted-quad notation (for example, 255.255.255.0).
ipv6-link-local	Specifies IPv6 and configures a link-local address generated by combining the standard link-local prefix with the automatically generated interface in the EUI-64 format. Using this option automatically generates an entire IPv6 address; this address is only a link-local, or VLAN-based, IPv6 address, that is, ports on the same segment can communicate using this IP address and do not have to pass through a gateway.
eui64	Specifies IPv6 and automatically generates the interface ID in the EUI-64 format using the interface's MAC address. Once you enter this parameter, you must add the following variables: <ipv6_address_mask>. Use this option when you want to enter the 64-bit prefix and use a EUI-64 address for the rest of the IPv6 address.
ipv6_address_mask	Specify the IPv6 address in the following format: x:x:x:x:x:x:prefix length, where each x is the hexadecimal value of one of the 8 16-bit pieces of the 128-bit wide address.

Default

N/A.

Usage Guidelines



NOTE

Use this command to assign an IP address to a vMAN, as well. The vMAN functionality is available only on platforms with an Advanced Edge, a Core, or an Advanced Core license.

The VLAN must already exist before you can assign an IP address: use the [create vlan](#) command to create the VLAN (also the vMAN must already exist).

**NOTE**

If you plan to use the VLAN as a control VLAN for an EAPS domain, do NOT configure the VLAN with an IP address.

**NOTE**

See [Chapter 22, “IP Unicast Commands,”](#) for information on adding secondary IP addresses to VLANs.

Beginning with ExtremeXOS software version 11.2, you can specify IPv6 addresses. See [Chapter 7, “IPv6 Unicast Routing,”](#) for information on IPv6 addresses.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS software version 11.3, you can assign an IP address (including IPv6 addresses) to a vMAN. Beginning with version 11.4, you can enable multicasting on that vMAN.

To enable multicasting on the specified vMAN once you assigned an IP address, take the following steps:

- 1 Enable IP multicast forwarding.
- 2 Enable and configure multicasting.

**NOTE**

You must upgrade to an Advanced Edge license to use vMAN functionality on a Summit X450e series switch.

BlackDiamond 10808 and 12800 series switches only. Beginning with ExtremeXOS software version 11.3, you can assign an IP address (including IPv6 addresses) to a specified vMAN to enable multicasting. You assign an IP address to a specified vMAN using the vMAN name, rather than a VLAN name.

To enable multicasting on the specified vMAN once you assigned an IP address, take the following steps:

- 1 Enable IP multicast forwarding.
- 2 Enable and configure multicasting.

**NOTE**

See [Chapter 29, “IP Multicast Commands,”](#) for information on configuring and using multicasting.

Example

The following commands are equivalent; both assign an IPv4 address of 10.12.123.1 to a VLAN named accounting:

```
configure vlan accounting ipaddress 10.12.123.1/24
configure vlan accounting ipaddress 10.12.123.1 255.255.255.0
```

The following command assigns a link local IPv6 address to a VLAN named management:

```
configure vlan accounting ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan name

```
configure vlan <vlan_name> name <name>
```

Description

Renames a previously configured VLAN.

Syntax Description

vlan_name	Specifies the current (old) VLAN name.
name	Specifies a new name for the VLAN.

Default

N/A.

Usage Guidelines

You cannot change the name of the default VLAN “Default.”



NOTE

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command renames VLAN *vlan1* to *engineering*:

```
configure vlan vlan1 name engineering
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan protocol

```
configure vlan <vlan_name> protocol <protocol_name>
```

Description

Configures a VLAN to use a specific protocol filter.

Syntax Description

vlan_name	Specifies a VLAN name.
protocol_name	Specifies a protocol filter name. This can be the name of a predefined protocol filter, or one you have defined. The following protocol filters are predefined: • IP • IPv6 • IPX • NetBIOS • DECNet • IPX_8022 • IPX_SNAP • AppleTalk any indicates that this VLAN should act as the default VLAN for its member ports.

Default

Protocol any.

Usage Guidelines

If the keyword `any` is specified, all packets that cannot be classified into another protocol-based VLAN are assigned to this VLAN as the default for its member ports.

Use the `configure protocol` command to define your own protocol filter.

Protocol filters on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. These devices do not forward packets with a protocol-based VLAN set to AppleTalk. To ensure that AppleTalk packets are forwarded on the device, create a protocol-based VLAN set to “any” and define other protocol-based VLANs for other traffic, such as IP traffic. The AppleTalk packets will pass on the “any” VLAN, and the other protocols will pass traffic on their specific protocol-based VLANs.

Example

The following command configures a VLAN named accounting as an IP protocol-based VLAN:

```
configure accounting protocol ip
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameter was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan tag

```
configure vlan <vlan_name> tag <tag> {remote-mirroring}
```

Description

Assigns a unique 802.1Q tag to the VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
tag	Specifies a value to use as an 802.1Q tag. The valid range is from 2 to 4095.
remote-mirroring	Specifies the tagged VLAN to be used for remote mirroring. NOTE: This variable is available only on the BlackDiamond 10808 and 12800 series switches.

Default

The default VLAN uses an 802.1Q tag (and an internal VLANid) of 1.

Usage Guidelines

If any of the ports in the VLAN will use an 802.1Q tag, a tag must be assigned to the VLAN. The valid range is from 2 to 4094 (tag 1 is assigned to the default VLAN, and tag 4095 is assigned to the management VLAN).

The 802.1Q tag will also be used as the internal VLANid by the switch.

You can specify a value that is currently used as an internal VLANid on another VLAN; it will become the VLANid for the VLAN you specify, and a new VLANid will be automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag (and internal VLANid) of 120 to a VLAN named *accounting*:

```
configure accounting tag 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create bvlan

```
create bvlan <bvlan-name>
```

Description

Configures the VLAN for the 802.1ah, or MAC-in-MAC tunneling, frames. Creates the 802.1ah backbone interface or tunnel entrance. Once you add the ports to this BVLAN, all ingress frames are encapsulated and all egress frames are stripped at these ports.

Syntax Description

bvlan-name	Enter the character string you want to name BVLAN.

Default

N/A.

Usage Guidelines

You can use any physical topology on the core backbone network. Although you can address IP addresses to backbone interfaces to test connectivity, do not enable IP forwarding.

You must use mutually exclusive names for:

- VLANs
- vMANs
- IPv6 tunnels
- BVLANs
- SVLANs

You can create and configure up to 4094 BVLANs on the switch.

Example

The following command creates the backbone VLAN (BVLAN) for the backbone interface:

```
create bvlan black
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create protocol

```
create protocol <name>
```

Description

Creates a user-defined protocol filter.

Syntax Description

name	Specifies a protocol filter name. The protocol filter name can have a maximum of 31 characters.

Usage Guidelines

Protocol-based VLANs enable you to define packet filters that the switch can use as the matching criteria to determine if a particular packet belongs to a particular VLAN.

After you create the protocol, you must configure it using the `configure protocol` command. To assign it to a VLAN, use the `configure vlan <vlan_name> protocol <protocol_name>` command.

Example

The following command creates a protocol named *fred*:

```
create protocol fred
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create svlan

```
create svlan <svlan-name>
```

Description

Creates the VLAN for the 802.1ad, or vMAN, frames. Each SVLAN carries the traffic for a different service provider, or service instance. Within the SVLAN, or vMAN, the customer VLANs remain untouched.

Syntax Description

svlan-name	Enter the character string you want to name SVLAN.

Default

N/A.

Usage Guidelines

This is a pure Layer 2 vMAN, which is transparent to users. Do not install Layer 2 control protocols on these interfaces because all BPDUs are tunneled through the 802.1ah, or MAC-in-MAC, tunnel. Do not assign any IP addresses to these interfaces, and ensure that IP forwarding is disabled.

You must use mutually exclusive names for:

- VLANs
- vMANs
- Ipv6 tunnels
- BVLANS
- SVLANs

You can create and configure up to 16 million SVLANs for each of the 4094 BVLANS supported on the switch, depending on the hardware resources of the switch.

Example

The following command creates the SVLAN white:

```
create svlan white
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

create vlan

```
create vlan <vlan_name> {vr <vr-name>}
```

Description

Creates a named VLAN.

Syntax Description

vlan_name	Specifies a VLAN name (up to 32 characters).
vr	Specifies a virtual router.
vr-name	Specifies in which virtual router to create the VLAN.
NOTE: On the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches, all VLANs are created in VR-Default and cannot be moved.	

Default

A VLAN named *Default* exists on all new or initialized Extreme switches:

- It initially contains all ports on a new or initialized switch, except for the management port(s), if there are any.
- It has an 802.1Q tag of 1.
- The default VLAN is untagged on all ports.
- It uses protocol filter *any*.

A VLAN named *Mgmt* exists on switches that have management modules or management ports:

- It initially contains the management port(s) the switch.
- It is assigned the next available internal VLANid as an 802.1Q tag.

If you do not specify the virtual router, the VLAN is created in the current virtual router.

Usage Guidelines

A newly-created VLAN has no member ports, is untagged, and uses protocol filter *any* until you configure it otherwise. Use the various `configure vlan` commands to configure the VLAN to your needs.

Internal VLANids are assigned automatically using the next available VLANid starting from the high end (4094) of the range.

Each VLAN name can be up to 32 standard alphanumeric characters, but must begin with an alphabetical letter. Quotation marks can be used to enclose a VLAN name that does not begin with an alphabetical character, or that contains a space, comma, or other special character.



NOTE

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

VLAN names are locally significant. That is, VLAN names used on one switch are only meaningful to that switch. If another switch is connected to it, the VLAN names have no significance to the other switch.

You must use mutually exclusive names for:

- VLANs
- vMANs
- Ipv6 tunnels
- BVLANs
- SVLANs

If you do not specify a virtual router when you create a VLAN, the system creates that VLAN in the default virtual router (VR-Default). The management VLAN is always in the management virtual router (VR-Mgmt).

Once you create virtual routers, ExtremeXOS software allows you to designate one of these as the domain in which all your subsequent configuration commands, including VLAN commands, are applied. If you create virtual routers, ensure that you are creating the VLANs in the desired virtual-router domain.

**NOTE**

The BlackDiamond 8800 series switches (formerly known as Aspen), SummitStack, and the Summit family of switches do not support user-created virtual routers; all user-created VLANs are in VR-Default.

For information on configuring virtual routers, see [Chapter 9, “Commands for Virtual Routers.”](#)

Example

The following command creates a VLAN named *accounting* on the current virtual router:

```
create vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The virtual router option `vr`, was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

delete bvlan

```
delete bvlan <bvlan-name>
```

Description

Deletes the specified VLAN for the 802.1ah, or MAC-in-MAC tunneling, frames, which is the 802.1ah backbone interface or tunnel entrance.

Syntax Description

bvlan-name	Enter the name of the BVLAN you are deleting.

Default

N/A.

Usage Guidelines

When you issue this command, the system automatically removes all SVLANs added to the specified BVLAN.

Example

The following command deletes the backbone VLAN (BVLAN) for the backbone interface:

```
delete bvlan black
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

delete protocol

```
delete protocol <name>
```

Description

Deletes a user-defined protocol.

Syntax Description

name	Specifies a protocol name.

Default

N/A.

Usage Guidelines

If you delete a protocol that is in use by a VLAN, the protocol associated with that VLAN will become none.

Example

The following command deletes a protocol named *fred*:

```
delete protocol fred
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete svlan

```
delete svlan <svlan-name>
```

Description

Deletes the SVLAN for the 802.1ad, or vMAN, frame. Each SVLAN carries the traffic for a different customer, or service instance.

Syntax Description

svlan-name	Enter the name of the SVLAN you are deleting.

Default

N/A.

Usage Guidelines

When you issue this command, the system automatically removes the specified SVLAN from its BVLAN, removes all ports added to the SVLAN, and removes the specified SVLAN from the system.

Example

The following command deletes the SVLAN white:

```
delete svlan white
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

delete vlan

```
delete vlan <vlan_name>
```

Description

Deletes a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If you delete a VLAN that has untagged port members and you want those ports to be returned to the default VLAN, you must add them back explicitly using the `configure svlan delete ports` command.



NOTE

The default VLAN cannot be deleted.

Example

The following command deletes the VLAN *accounting*:

```
delete accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable dot1p examination inner-tag ports

```
disable dot1p examination inner-tag ports [all | <port_list>]
```

Description

Used with vMANs, and instructs the switch to examine the 802.1p value of the outer tag, or added vMAN header, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
port_list	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the outer tag, or vMAN encapsulation tag, when assigning the packet to an egress queue at the egress port of the vMAN.



NOTE

See [Chapter 12, "QoS Commands,"](#) for information on configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value.

Example

The following command uses the 802.1p value on the outer tag, or vMAN encapsulation, to put the packet in the egress queue on the vMAN egress port:

```
disable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

disable loopback-mode vlan

```
disable loopback-mode vlan <vlan_name>
```

Description

Disallows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command disallows the VLAN *accounting* to be placed in the UP state without an external active port:

```
disable loopback-mode vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable vlan

```
disable vlan <vlan-name>
```

Description

Use this command to disable the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN you want to disable.

Default

Enabled.

Usage Guidelines

This command allows you to administratively disable specified VLANs. The following guidelines apply to working with disabling VLANs:

- Disabling a VLAN stops *all* traffic on all ports associated with the specified VLAN.
- You cannot disable any VLAN that is running any Layer 2 protocol traffic at all.

When you attempt to disable a VLAN running Layer 2 protocol traffic, the system returns a message similar to the following:

```
VLAN accounting cannot be disabled because it is actively use by an L2 Protocol
```

- You can disable the default VLAN; ensure that this is necessary prior to disabling the default VLAN.
- You *cannot* disable the management VLAN.
- Although you can remove ports from a disabled VLAN, you *cannot* add ports to a disabled VLAN or bind Layer 2 protocols to that VLAN.

When you attempt to disable a VLAN running Layer 2 protocol traffic, the system returns a message similar to the following:

```
VLAN accounting is disabled. Enable VLAN before adding ports.
```

Example

The following command disables the VLAN named accounting:

```
disable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable dot1p examination inner-tag port

```
enable dot1p examination inner-tag port [all | <port_list>]
```

Description

Used with vMANs, and instructs the switch to examine the 802.1p value of the inner tag, or header of the original packet, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
port_list	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the inner, or original, tag when assigning the packet to an egress queue at the egress port of the vMAN.



NOTE

See [Chapter 12, "QoS Commands,"](#) for information on configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value.

Example

The following command puts the packets in the egress queue of the vMAN egress port according to the 802.1p value on the inner tag:

```
enable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

enable loopback-mode vlan

```
enable loopback-mode vlan <vlan_name>
```

Description

Allows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLAN's routing interface to become active.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command allows the VLAN *accounting* to be placed in the UP state without an external active port:

```
enable loopback-mode vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable vlan

```
enable vlan <vlan-name>
```

Description

Use this command to re-enable a VLAN that you previously disabled.

Syntax Description

vlan_name	Specifies the VLAN you want to disable.

Default

Enabled.

Usage Guidelines

This command allows you to administratively enable specified VLANs that you previously disabled.

Example

The following command enables the VLAN named accounting:

```
enable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show bvlan

```
show bvlan {<bvlan_name> {ipv4 | ipv6}} {detail}
```

Description

Displays BVLAN configuration. BVLANS are used in MAC-in-MAC tunneling (IEEE 802.1ah) configurations.

Syntax Description

bvlan_name	Specifies information is displayed for particular BVLAN.
ipv4	Not supported in ExtremeXOS version 11.4.
ipv6	Not supported in ExtremeXOS version 11.4.
detail	Specifies all information is displayed for each BVLAN.

Default

Summary information for all BVLANS on the device.

Usage Guidelines

None.



NOTE

To display the EtherType for IEEE 802.1ah MAC-in-MAC tunneling (BVLAN), use the command `show vman ethertype`.

Example

The following is an example of the display from the `show bvlan` command:

```
* BD-10808.6 # show bvlan
```

```
-----
Name                VID  Protocol Addr          Flags                Proto  Ports  Virtual
                        Active router
                        /Total
-----
test                4091 -----B      ANY      0 /0  VR-Default
-----
```

```
Flags : (C) EAPS Control vlan, (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback
        Enabled, (m) IPmc Forwarding Enabled,
        (n) IP Multinetting Enabled, (N) Network LogIn vlan,
        (o) OSPF Enabled, (p) PIM Enabled,
        (P) EAPS protected vlan, (r) RIP Enabled, (T) Member of STP Domain,
        (v) VRRP Enabled, (B) 802.1ah Backbone VMAN, (S) 802.1ah Service VMAN
Total number of bvlan(s) : 1
```

The following is an example of the display from the `show bvlan <bvlan_name>` command:

```
BVLAN Interface with name test created by user
  Admin State:      Enabled      Tagging:Untagged (Internal tag 4091)
  Priority:          802.1P Priority 0
  Virtual router:   VR-Default
  IPv6:             None
  STPD:             None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disabled
  NetLogin:         Disabled
  QosProfile:       QP1
  Dot1ah Mode:      Backbone
  Service Count:    0
  Ports:           0.           (Number of active ports=0)
```

The display from the `show bvlan detail` command shows all the information shown in the `show bvlan <bvlan_name>` command, but displays information for all configured BVLANs.

Displaying EtherTypes. The following is an example of the display from the `show vman etherType` command when you configure MAC-in-MAC tunnels:

```
BlackDiamond 12804.41 # show vman etherType
vman EtherType : 0x88a8
bvlan EtherType: 0x88b5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show mac-binding

```
show mac-binding {bvlan <bvlan-name> {<bridge-mac> {svlan <svlan-name>  
<customer-mac>}} | svlan <svlan-name>}
```

Description

Displays information about the internal database binding the remote bridge (BVLAN) MAC addresses and customer host MAC addresses behind that BVLAN.

Syntax Description

bvlan-name	Enter the name of the BVLAN you want to display.
bridge-mac	Enter the MAC address of the BVLAN you want to display.
svlan-name	Enter the name of the SVLAN you want to display.
customer-mac	Enter the name of the customer-mac you want to display.

Default

N/A.

Usage Guidelines

Use this command to display information on the database binding the backbone and customer MAC addresses.

This command displays the following information:

- BVLAN name
- B-TAG value
- MAC address of the 802.1ah-configured bridge facing the customer's device
- SVLAN
- S-TAG value
- MAC address of device that sent the frame
- Age of the entry in seconds
- Port number on which the packet was learned

Example

The following command displays the MAC binding table on the switch:

```
show mac-binding
```

The following is sample output from this command:

```
# BD-10808.1 # show mac-binding
BVLAN VID Mac                SVLAN VID MAC                Flags    AGE    Port
B      100 00:01:30:F9:9E:60   S1      10 00:A0:00:00:00:01      d       60    2:1
B      100 00:01:30:F9:9E:61   S2      20 00:B0:00:00:00:02      d
```

Flags : d - Dynamic, s - Static, p - Permanent

Total: 0 Static: 0 Perm: 2 Dyn: 0

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show protocol

```
show protocol {<name>}
```

Description

Displays protocol filter definitions.

Syntax Description

name	Specifies a protocol filter name.

Default

Displays all protocol filters.

Usage Guidelines

Displays the defined protocol filter(s) with the types and values of its component protocols.

Example

The following is an example of the `show protocol` command:

Protocol Name	Type	Value
IP	etype	0x0800
	etype	0x0806
ANY	ANY	0xffff
ipx	etype	0x8137
decnet	etype	0x6003
	etype	0x6004
netbios	llc	0xf0f0
	llc	0xf0f1
ipx_8022	llc	0xe0e0
ipx_snap	snap	0x8137
appletalk	snap	0x809b
	snap	0x80f3

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show svlan

```
show svlan {<svlan_name> {ipv4 | ipv6}} {detail}
```

Description

Displays SVLAN configuration. SVLANs are used in MAC-in-MAC tunneling (IEEE 802.1ah) configurations.

Syntax Description

svlan_name	Specifies information is displayed for particular SVLAN.
ipv4	Not supported in ExtremeXOS version 11.4.
ipv6	Not supported in ExtremeXOS version 11.4.
detail	Specifies all information is displayed for each SVLAN.

Default

Summary information for all SVLANs on the device.

Usage Guidelines

None.



NOTE

To display the EtherType for IEEE 802.1ah MAC-in-MAC tunneling, use the command `show vman ethertype`.

Example

The following is an example of the display from the `show svlan` command:

```
* BD-10808.8 # show svlan
```

Name	VID	Protocol	Addr	Flags	Proto	Ports Active router /Total	Virtual
servtest	4090	-----	-----	-----S	ANY	0 /0	VR-Default

```
Flags : (C) EAPS Control vlan, (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback
        Enabled, (m) IPmc Forwarding Enabled,
        (n) IP Multinetting Enabled, (N) Network LogIn vlan,
        (o) OSPF Enabled, (p) PIM Enabled,
        (P) EAPS protected vlan, (r) RIP Enabled, (T) Member of STP Domain,
        (v) VRRP Enabled, (B) 802.1ah Backbone VMAN, (S) 802.1ah Service VMAN
Total number of svlan(s) : 1
```

**NOTE**

The show svlan <svlan_name> display may differ slightly depending on the platform and configuration you are using.

The following is sample output for the show svlan <svlan_name> command when you are displaying a service VLAN (SVLAN) for a MAC-in-MAC tunnel:

```
BlackDiamond 12804.36 # show svlan servtest
SVLAN Interface with name servtest created by user
  Admin State:      Enabled          Tagging:      802.1Q Tag 200
  Priority:         802.1P Priority 0
  Virtual router:   VR-Default
  IPv6:            None
  STPD:            None
  Protocol:        Match all unfiltered protocols
  Loopback:        Disable
  NetLogin:        Disabled
  QosProfile:       QP1
  Dot1ah Mode:     Service
  I-SID:           250
  In Backbone:     backbone
  Ports:           0.              (Number of active ports=0)
```

**NOTE**

If you did not configure an ISID, the I-SID line does not display. (The system uses the value of the S-tag if you do not optionally configure an ISID.)

The display from the show svlan detail command shows all the information shown in the show svlan <svlan_name> command, but displays information for all configured SVLANs.

Displaying EtherTypes. The following is an example of the display from the show vman etherType command when you configure MAC-in-MAC tunnels (802.1ah):

```
BlackDiamond 12804.41 # show vman etherType
vman EtherType : 0x88a8
bvlan EtherType: 0x88b5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 10808 and 12800 series switches.

show vlan

```
show vlan {detail {ipv4 | ipv6} | <vlan_name> {ipv4 | ipv6} | virtual-
router <vr-router> | <vlan_name> stpd | security}
```

Description

Displays information about VLANs.

Syntax Description

detail	Specifies that detailed information should be displayed for each VLAN.
vlan_name	Specifies a VLAN name.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.
vr-name	Specifies a virtual router name. NOTE: This variable is available only on the BlackDiamond 10808 and 12800 series switches.
stpd	Specifies that STP domains displays for each VLAN.
security	Enables security checking

Default

Summary information for all VLANs on the device.

Usage Guidelines



NOTE

To display IPv6 information, you must issue either the `show vlan detail` command or `show vlan` command with the name of the specified VLAN.

Unlike many other VLAN-related commands, the keyword “vlan” is required in all forms of this command except when requesting information for a specific vlan.

Use the command `show vlan` to display summary information for all VLANs. It shows various configuration options as a series of “flags” (see the example below). VLAN and protocol names may be abbreviated in this display.

Use the command `show vlan detail` to display detailed information for all VLANs. This displays the same information as for an individual VLAN, but shows every VLAN, one-by-one. After each VLAN display you can elect to continue or quit.

Protocol `none` indicates that this VLAN was configured with a user-defined protocol that has subsequently been deleted.

**NOTE**

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches display the Mgmt VLAN in the VR-Mgmt; they display all other VLANs in the VR-Default.

When an IPv6 address is configured for the VLAN, the system may display one of the following two address types in parentheses after the IPv6 address:

- Tentative
- Duplicate

**NOTE**

See the ExtremeXOS Concepts Guide for information on IPv6 address types.

You can display additional useful information on VLANs configured with IPv6 addresses by issuing the `show ipconfig ipv6 vlan <vlan_name>`. The following is sample output from this command:

```
BD10K # show ipconfig ipv6 my_ipv6_100
Router Interface on my_ipv6_100 is enabled and up. MTU: 1500
Locally registered unicast addresses:
2001:db8::8:802:200c:417a/64
fe80::230:48ff:fe41:ed97%my_ipv6_100/64
Flags:
IPv6 Forwarding: YES Accept recvd RA: NO
Send redirects: NO Accept redirects: NO
```

Example

The following is an example of the `show vlan` command on the BlackDiamond 10808 switch:

```
-----
Name                VID Protocol Addr          Flags          Proto Ports  Virtual
                  Active router
                  /Total
-----
Default             1  -----T--- ANY    0 /771 VR-Default
hugheapscontrol     4092 -----C ANY    0 /0   VR-Default
Mgmt                4095 ----- ANY    1 /1   VR-Mgmt
peggy               4094 10.0.0.1 /8 -L----- ANY    1 /2   VR-Default
santaclara          4091 -----D-----d ANY    0 /1   benton
-----
Flags : (C) EAPS Control vlan, (d) NetLogin Dynamically created VLAN,
        (D) VLAN Admin Disabled, (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enab
led,
        (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
        (N) Network LogIn vlan, (o) OSPF Enabled, (p) PIM Enabled,
        (P) EAPS protected vlan, (r) RIP Enabled,
        (T) Member of STP Domain, (v) VRRP Enabled

Total number of Vlan(s) : 5
```

The following is an example of the `show vlan Default` command:

```
VLAN Interface with name Default created by user
  Admin State:      Enabled      Tagging:      802.1Q Tag 1
  Priority:          802.1P Priority 0
  Virtual router:   VR-Default
  Primary IP:       10.0.0.1/8
  Secondary IPs:    12.0.0.5/8
  IPv6:             None
  STPD:             s0(Disabled,Auto-bind)
  Protocol:         Match all unfiltered protocols
  Loopback:         Disable
  NetLogIn:         Enabled
  QosProfile:       QP1
  Ports: 771.       (Number of active ports=0)
    Untag:          3:2,      4:1,      4:2,      5:1,      5:2,      5:3,      5:4,
                   5:5,      5:6,      5:7,      5:8,      5:9,      5:10,     5:11,
                   5:12,     5:13,     5:14,     5:15,     5:16,     5:17,     5:18,
                   5:19,     5:20,     5:21,     5:22,     5:23,     5:24,     5:25,
                   5:26,     5:27,     5:28,     5:29,     5:30,     5:31,     5:32,
                   5:33,     5:34,     5:35,     5:36,     5:37,     5:38,     5:39,
                   5:40,     5:41,     5:42,     5:43,     5:44,     5:45,     5:46,
                   5:47,     5:48,     5:49,     5:50,     5:51,     5:52,     5:53,
                   5:54,     5:55,     5:56,     5:57,     5:58,     5:59,     5:60,
                   5:61,     5:62,     5:63,     5:64,     5:65,     5:66,     5:67,
                   5:68,     5:69,     5:70,     5:71,     5:72,     5:73,     5:74,
                   5:75,     5:76,     5:77,     5:78,     5:79,     5:80,     5:81,
                   5:82,     5:83,     5:84,     5:85,     5:86,     5:87,     5:88,
  Flags:            (*) Active, (!) Disabled, (g) Load Sharing port
                   (b) Port blocked on the vlan, (a) Authenticated NetLogin port
                   (u) Unauthenticated NetLogin port, (m) Mac-Based port
```



NOTE

The `m` flag for MAC-based port represents network login information; it is only supported on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.



NOTE

The `number of active ports` line displays the number of ports presently in forwarding state on this VLAN.

The output for the `show vlan detail` command displays the same information for all VLANs configured on the switch.



NOTE

See [Chapter 22, "IP Unicast Commands,"](#) for information on adding secondary IP addresses to VLANs.

History

This command was first available in ExtremeXOS 10.1.

The IPv6 information was added in ExtremeXOS 11.2.

The netlogin information was added in ExtremeXOS 11.3.

The virtual router and administratively enabled/disabled information was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

Information on MAC-based ports is available only on the Summit family of switches, SummitStack, and the BlackDiamond 8800 series switch.

Information on the virtual routers is available only on the BlackDiamond 10808 and 12800 series switches.

unconfigure vlan ipaddress

```
unconfigure vlan <vlan_name> ipaddress {<ipv6_address_mask>}
```

Description

Removes the IP address of the VLAN or a vMAN. With no parameters, the command removes the primary IPv4 address on the specified VLAN. Using the IPv6 parameters, you can remove specified IPv6 addresses from the specified VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
ipv6_address_mask	Specifies an IPv6 address using the format of IPv6-address/prefix-length, where IPv6 is the 128-bit address and the prefix length specifies the number of leftmost bits that comprise the prefix.

Default

Removes the primary IPv4 address from the specified VLAN.

Usage Guidelines



NOTE

You need an Advanced Edge, a Core, or an Advanced Core license to use vMANs.

If you do not specify any parameters, this command removes the primary IPv4 address from the VLAN.



NOTE

With IPv6, you cannot remove the last link local IPv6 address until all global IPv6 addresses are removed.

Beginning with ExtremeXOS software version 11.3, you can use this command to unconfigure the IP address for a vMAN as well.

Example

The following command removes the primary IPv4 address from the VLAN accounting:

```
unconfigure vlan accounting ipaddress
```

The following command removes an IPv6 addresses from the VLAN finance:

```
unconfigure vlan finance ipaddress 3ffe::1
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

13 VMAN Commands

This chapter describes commands for creating, deleting, and performing basic performing basic virtual metropolitan area network vMAN™ configurations.

configure flood-group extended add

```
configure flood-group extended <fg_name> add <vman_name> port <port_list |
all>
```

Description

Adds one or more ports to an extended flood group in a vMAN.

Syntax Description

fg_name	Enter the name of the extended flood group.
vman_name	Enter the vMAN name.
port_list	Enter the list of ports to add to the vMAN.
all	Adds all the ports in the vMAN.

Default

N/A.

Usage Guidelines

Use this feature to add one or more ports to an extended flood group.

This feature allows you to flood packets to a crafted set of ports, known as a flood group. The ports can be a member of any vMANs in the switch. These flood groups can be crafted to suit your requirements and be applied to a vMAN. This enables packets from the vMAN to be flooded based on the flood groups, allowing a single packet to get distributed to *n* ports.



NOTE

Configure the trunk ports before adding ports to or deleting ports from the flood group.

Example

The following command adds port 1:1 to extended flood group *fred* for vMAN *accounting*:

```
configure flood-group extended fred add accounting port 1:1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

configure flood-group extended delete

```
configure flood-group extended <fg_name> delete <vman_name> port <port_list>
| all>
```

Description

Deletes one or more ports from an extended flood group in a vMAN.

Syntax Description

fg_name	Enter the name of the extended flood group.
vman_name	Enter the vMAN name.
port_list	Enter the list of ports to delete from the vMAN.
all	Deletes all the ports from the vMAN.

Default

N/A.

Usage Guidelines

Use this feature to delete one or more ports from an extended flood group.

This feature allows you to flood packets to a crafted set of ports, known as a flood group. The ports can be a member of any vMANs in the switch. These flood groups can be crafted to suit your requirements and be applied to a vMAN. This enables packets from the vMAN to be flooded based on the flood groups, allowing a single packet to get distributed to *n* ports.



NOTE

Configure the trunk ports before adding ports to or deleting ports from the flood group.

Example

The following command deletes port 1:1 to extended flood group fred for vMAN *accounting*:

```
configure flood-group extended fred delete accounting port 1:1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

configure flood-group virtual add

```
configure flood-group virtual <fg_name> add <vman_name> stag <tag> port
<port_list | all>
```

Description

Adds one or more ports to a virtual flood group in a vMAN.

Syntax Description

fg_name	Enter the name of the extended flood group.
vman_name	Enter the vMAN name.
tag	Enter the value of the s-ethertype you want to use.
port_list	Enter the list of ports to add to the vMAN flood group.
all	Adds all the ports to the vMAN flood group.

Default

N/A.

Usage Guidelines

Use this command to add one or more ports to a virtual flood group.

This feature allows you to flood packets to a crafted set of ports, known as a flood group, with crafted stag. The ports can be a member of any vMANs in the switch. Each port in the virtual flood-group has a stag that does not belong to any other VLAN or VMAN. These flood groups can be crafted to suit your requirements and be applied to a vMAN. This enables packets from the vMAN to be flooded based on the flood groups, allowing a single packet to get distributed to *n* ports.



NOTE

Configure the trunk ports before adding ports to or deleting ports from the flood group.

Example

The following command adds port 1:1 to virtual flood group *fred* for vMAN *accounting*:

```
configure flood-group virtual fred add accounting port 1:1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

configure flood-group virtual delete

```
configure flood-group virtual <fg_name> delete <vman_name> stag <tag> port
<port_list | all>
```

Description

Deletes one or more ports from a virtual flood group in a vMAN.

Syntax Description

fg_name	Enter the name of the extended flood group.
vman_name	Enter the vMAN name.
tag	Enter the value of the s-ethertype you want to use.
port_list	Enter the list of ports to delete from the vMAN flood group.
all	Deletes all the ports from the vMAN flood group.

Default

N/A.

Usage Guidelines

Use this command to delete one or more ports from a virtual flood group.

This feature allows you to flood packets to a crafted set of ports, known as a flood group. The ports can be a member of any vMANs in the switch. These flood groups can be crafted to suit your requirements and be applied to a vMAN. This enables packets from the vMAN to be flooded based on the flood groups, allowing a single packet to get distributed to *n* ports.



NOTE

Configure the trunk ports before adding ports to or deleting ports from the flood group.

Example

The following command deletes port 1:1 to virtual flood group fred for vMAN *accounting*:

```
configure flood-group virtual fred delete accounting port 1:1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

configure port ethertype

```
configure port <port_list> ethertype {primary | secondary}
```

Description

Assigns a port to the primary or secondary Ethertype.

Syntax Description

port_list	Specifies the list of ports to be configured.
primary	Specifies that the Ethertype is primary.
secondary	Specifies that the Ethertype is secondary.

Default

N/A.

Usage Guidelines

Use this command to assign a port or list of ports to the primary or secondary vMAN Ethertype.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

create flood-group extended

```
create flood-group extended <fg_name>
```

Description

Creates an extended flood group.

Syntax Description

fg_name	Enter the name of the extended flood group you want to create.

Default

N/A.

Usage Guidelines

Use this command to create a virtual flood list to allow you to flood to *n* vMANs. This enables packets from the vMAN to be flooded based on these flood lists. Thus a single packet gets distributed to *n* ports.

Example

The following command creates a extended flood group with the name *fred*:

```
create flood-group extended fred
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

create flood-group virtual

```
create flood-group virtual <fg_name>
```

Description

Creates a virtual flood group.

Syntax Description

fg_name	Enter the name of the virtual flood group you want to create.

Default

N/A.

Usage Guidelines

Use this command to create a virtual flood list to allow you to flood to *n* vMANs. This enables packets from the vMAN to be flooded based on these flood lists. Thus a single packet gets distributed to *n* ports.

Note the following with regard to the flood group policy file keyword:

- at least one *fg_name* must present
- a flood group may have no port
- the flood group keyword is valid only in ingress ACL
- enabling or disabling of the sharing of the trunk-port is not reflected to the flood group

Example

The following command creates a virtual flood group with the name *fred*:

```
create flood-group virtual fred
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the BlackDiamond 12800 series switches only.

configure vman add ports

```
configure vman <vman-name> add ports [ all | <port_list> ] {untagged |
tagged | nobroadcast }
```

Description

Adds one or more ports in a vMAN.

Syntax Description

vman-name	Specifies a vMAN name.
all	Specifies all ports in the vMAN.
port_list	Specifies a list of ports or slots and ports.
untagged	Specifies that the ports are untagged.
tagged	Specifies that the ports are tagged.
nobroadcast	Specifies that the ports do not forward broadcast, multicast, or unknown unicast packets. NOTE: This option is available only on the BlackDiamond 10808 and 12800 series switch.

Default

N/A.

Usage Guidelines

If you do not specify a parameter, the default value is untagged.

The vMAN tunnel begins at the ingress, or customer access, port and terminates at the egress, or trunk, port. Traffic flows from the egress trunk port onto the network thereafter without the vMAN tag. Ensure that all the switch-to-switch ports in the vMAN tunnel are configured as tagged ports. Configure the vMAN ingress, or customer access, port as an untagged port (although this port does accept tagged packets). You must configure the vMAN tunnel egress, or trunk, port as an untagged port so that the vMAN header is stripped from the frame.



NOTE

You must configure the vMAN tunnel egress, or trunk, port as untagged so that the vMAN header is stripped from the frame.

The vMAN must already exist before you can add (or delete) ports: use the `create vman` command to create the VLAN. vMAN ports can belong to load-sharing groups.

The Extreme Networks default Ethernet type for vMAN is 088a8.

You cannot configure Connectivity Fault Management (CFM) on ports added to the vMAN.

**NOTE**

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS 11.6, you can configure vMANs and *untagged* VLANs on the same physical port regardless of the Ethernet type for the vMAN on any BlackDiamond 8800 switches, SummitStack, and the Summit family of switches.

Beginning with ExtremeXOS 11.5, you can configure vMANs and VLANs on the same module or on the same stand-alone switch for the following devices: BlackDiamond 8800 a-series and e-series modules, Summit X450a and X450e series switches, and the SummitStack. You can configure both vMAN and VLAN simultaneously on the same physical port, but you must first change the vMAN Ethernet type to 0x8100.

Beginning with ExtremeXOS 11.4, you can configure vMANs and VLANs on the same module, on the same stand-alone switch, or on the same port once you configure the vMAN Ethernet type as 0x8100. You must specifically configure this vMAN Ethernet type; the default vMAN Ethernet type is 0x88a8. If you attempt to configure vMANs and VLANs on the same module or stand-alone switch without changing the Ethernet type to 0x8100, the system returns an error.

You must enable jumbo frames for the entire module or stand-alone switch prior to creating and configuring vMANs. Use the `enable jumbo-frame ports all` command to enable jumbo frames on the BlackDiamond 8800 switches, SummitStack, and Summit switches.

BlackDiamond 10808 and 12800 series switch only. A physical port configured in a vMAN as an *untagged* port *cannot* also belong to a VLAN; however, a port configured in a vMAN as a *tagged* port *can* simultaneously belong to a VLAN.

If you specify the parameter `nobroadcast`, the system used the default value of `untagged`; you cannot add `nobroadcast` once you have specified tagged packets.

All vMAN ports are automatically enabled for jumbo frames to accommodate the extra vMAN tag. If any port in the load-sharing group is enabled for vMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, vMAN is automatically enabled on all ports of the untagged load-sharing group.

All ports added to a specified vMAN must be in the same virtual router. For more information on displaying, configuring, and using virtual routers, see [Chapter 9, "Commands for Virtual Routers."](#)

Example

The following command assigns ports 1:1, 1:2, 1:3, and 1:6 to a vMAN named *accounting*:

```
configure vman accounting add ports 1:1, 1:2, 1:3, 1:6 tag 100
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure vman delete ports

```
configure vman <vman-name> delete ports [ all | <port_list> ]
```

Description

Deletes one or more ports in a vMAN.

Syntax Description

vlan_name	Specifies a vMAN name.
all	Specifies all ports in the vMAN.
port_list	Specifies a list of ports or slots and ports.

Default

N/A.

Usage Guidelines

The vMAN must already exist before you can add (or delete) ports: use the `create vman` command to create the VLAN.

Example

The following command deletes ports 1:1, 1:2, 1:3, and 1:6 on a modular switch to a vMAN named *accounting*:

```
configure vman accounting delete ports 1:1, 1:2, 1:3, 1:6
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

configure vman ethertype

```
configure vman ethertype <value> [ primary | secondary ]
```

Description

Changes the default Ethernet type for the vMAN header; the default Extreme Networks Ethernet type is 0x88A8.

Syntax Description

value	Specifies an Ethertype value in the format of 0xffff.
primary	Specifies that the Ethertype is primary.
secondary	Specifies that the Ethertype is secondary.

Default

Ethertype value of 0x88a8 and type primary.

Usage Guidelines

The default vMAN Ethernet type on Extreme devices is 0x88a8.

If your vMAN transits a third-party device (other than an Extreme Networks device), you must configure the EtherType for the vMAN tag as the Ethernet type that the third-party device uses.

Beginning with ExtremeXOS 11.3, the system supports all vMAN EtherTypes, including the standard Ethernet type of 0x8100. Beginning with ExtremeXOS 11.6, you can configure vMANs and *untagged* VLANs on the same physical *port* regardless of the Ethernet type for the vMAN on any BlackDiamond 8800 switches, SummitStack, and the Summit family of switches.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS 11.5, you can configure vMANs and VLANs on the same module or on the same stand-alone switch for the following devices: BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches, whether or not included in a SummitStack. You can configure both vMAN and VLAN simultaneously on the same physical port, but you must first change the vMAN Ethernet type to 0x8100.

Beginning with ExtremeXOS 11.4, you can configure vMANs and VLANs on the same module, on the same stand-alone switch, or the same port once you configure the vMAN Ethernet type as 0x8100. You must specifically configure this vMAN Ethernet type; the default vMAN Ethernet type is 0x88a8. If you attempt to configure vMANs and VLANs on the same module or stand-alone box without changing the Ethernet type to 0x8100, the system returns an error.

BlackDiamond 10808 and 12800 series switches only. A physical port configured in a vMAN as an *untagged* port *cannot* also belong to a VLAN; however, a port configured in a vMAN as a *tagged* port *can* simultaneously belong to a VLAN.

Example

The following command, followed by a switch reboot, changes the Ethertype value to 8100:

```
configure vman ethertype 0x8100
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license. For BlackDiamond 12800 series switches only, the secondary keyword for Ethertype is supported; by default this value is 0.

configure vman tag

```
configure vman <vman-name> tag <tag>
```

Description

Assigns a unique 802.1Q tag to the vMAN.

Syntax Description

vlan_name	Specifies a vMAN name.
tag	Specifies a value to use as vMAN tag 802.1Q tag. The valid range is from 2 to 4094.

Default

N/A.

Usage Guidelines

You add a tag value to each vMAN you configure. This is then carried in the 4-byte vMAN header, which is added to all packets (both tagged and untagged) that arrive on the port you added to the vMAN.

You can specify a value that is currently used as an internal VLANid on another VLAN; it will become the VLANid for the VLAN you specify, and a new VLANid will be automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag of 120 to a vMAN named *accounting*:

```
configure vman accounting tag 120
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

create vman

```
create vman <vman-name> {vr <vr_name>}
```

Description

Creates a vMAN, which allows you to create a “tunnel” through the public network.

Syntax Description

vman-name	Specifies a vMAN name. Up to 32 characters.
vr	Specifies a virtual router.
vr_name	Specifies the virtual router. NOTE: On the BlackDiamond 8800 series switches, SummitStack, and the Summit family switches, all vMANs are created in VR-Default and cannot be moved.

Default

N/A

Usage Guidelines

vMANs allow you to establish a private path through the public network.

After you create the vMAN, you must configure the vMAN tag plus adding the ports that you want.



NOTE

On the BlackDiamond 10808 and 12800 series switch, all vMAN ports are automatically enabled for jumbo frames to accommodate the added vMAN tag. On the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches, you must enable jumbo frames on the switch.

If you do not specify the virtual router, the vMAN is created in the current virtual router.



NOTE

Beginning with ExtremeXOS 11.6, you can configure vMANs and untagged VLANs on the same physical port, regardless of other Ethernet type for the vMANs. With ExtremeXOS 11.5, you can configure VLANs and vMANs on the same module on the BlackDiamond 8800 a-series and e-series modules and on the same stand-alone Summit X450a and X405e series switch, whether or not included in a SummitStack. With ExtremeXOS 11.4, to configure both VLANs and vMANs on the same module on the BlackDiamond 8800 series switches, SummitStack, and on the Summit family of switches, you must change the vMAN Ethernet type to 8100. See [configure vman ethertype](#) to change the Ethernet type for vMANs.

To configure a vMAN, follow these steps:

- 1 Create the tunnel by creating the vMAN.
- 2 If you are working on the BlackDiamond 8800 switches, SummitStack, or the Summit switches, enable jumbo frames.
Use the following command:
`enable jumbo-frames ports all`
- 3 Assign a tag value to the vMAN.
Use the following command:
`configure vman <vman-name> tag <tag>`
- 4 Add the ports in the tunnel to the vMAN.
- 5 Configure vMAN member ports as tagged on switch-to-switch ports and untagged on the ingress and egress ports of the tunnel.

Use the following command:

```
configure vman <vman-name> add ports [ all | <port_list> ] {untagged | tagged | nobroadcast }
```



NOTE

You must configure the vMAN tunnel egress, or trunk, port as untagged so that the vMAN header is stripped from the frame.

The Extreme Networks default Ethernet type for vMAN is 088a8.

You must use mutually exclusive names for:

- VLANs
- vMANs
- IPv6 tunnels
- SVLANs
- BVLANs



NOTE

Beginning with ExtremeXOS software version 11.3 you can tunnel multicast traffic on the BlackDiamond 10808 and 12800 series switch. To assign an IP address to a vMAN, use the `configure vlan ipaddress` command.

Example

The following command creates a vMAN named *fred*:

```
create vman fred
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

delete vman

```
delete vman <vman-name>
```

Description

Deletes a previously created vMAN.

Syntax Description

vman-name	Specifies a vMAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the vMAN *accounting*:

```
delete vman accounting
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

show vman

```
show vman {<vlan_name> {eaps | ipv4 | ipv6} | etherType | detail {ipv4 |
ipv6}}
```

Description

Displays vMAN information including configuration and EtherType.



NOTE

The information displayed on this command depends on the platform and configuration you are using.

Syntax Description

vlan_name	Specifies information is displayed for particular vMAN.
eaps	Displays the EAPS domains to which the vMAN belongs.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.
etherType	Specifies that EtherType information is vMAN and, if configured, BVLAN (MAC-in-MAC) applications.
detail	Specifies all information is displayed for each vMAN.

Default

Summary information for all vMANs on the device.

Usage Guidelines

The information displayed with this command depends on the platform and configuration you are using.

Display on the BlackDiamond 10808 and 12800 series switch only. If you configured MAC-in-MAC tunneling (802.1ah) on the switch, the `show vman` displays information on the BVLANS and the SVLANS as well. (You can display the same information using the `show bvlan` or `show svlan` command.



NOTE

To display the Ether Type for MAC-in-MAC tunneling (802.1ah), you must use the `show vman etherType` command.

Display on the BlackDiamond 12800 R-series switch only. If you configured vMAN ACLs on the BlackDiamond 12800 R-series switch and you designated one port to handle flooding, broadcast, and multicast traffic egressing the vMAN, use the `show vman <vman_name> detail` command to display the information.

Example

The following is an example of the display from the `show vman` command:

```
-----
Name                VID Protocol Addr          Flags          Proto Ports Virtual
                  Active router
                  /Total
-----
peggy                100  10.10.255.3    /24    f-----m----- ANY    0 /0    VR-Default
-----
Flags : (C) EAPS Control vlan, (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback
        Enabled, (m) IPmc Forwarding Enabled,
        (n) IP Multinetting Enabled, (N) Network LogIn vlan,
        (o) OSPF Enabled, (p) PIM Enabled,
        (P) EAPS protected vlan, (r) RIP Enabled, (T) Member of STP Domain,
        (v) VRRP Enabled, (B) 802.1ah Backbone VMAN, (S) 802.1ah Service VMAN
```

Total number of vman(s) : 1

The following is an example of the display from the `show vman <vlan_name>` command:

```
VLAN Interface with name test created by user
  Admin State: Enabled          Tagging:Untagged (Internal tag 4090)
  Priority:          802.1P Priority 0
  Virtual router: VR-Default
  Primary IP       : 10.10.255.3/24
  IPv6:            NONE
  STPD:            None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disable
  NetLogIn:         Enabled
  Rate Shape:       Disable
  QosProfile:       QP1
  Ports: 2.         (Number of active ports=2)
    Flags:          (*) Active, (!) Disabled
                  (g) Load Sharing port
    Untag:          *3:1
    Tag:            *3:2
```

The display from the `show vman detail` command shows all the information shown in the `show vman <vlan_name>` command, but displays information for all configured vMANs.

Displaying EtherTypes. The following is an example of the display from the `show vman etherType` command:

```
vMan EtherType: 0x88a8
```

The following is an example of the display from the `show vman etherType` command when you configure MAC-in-MAC tunnels (BlackDiamond 10808 and 12800 series switches only):

```
BlackDiamond 12804.41 # show vman etherType
vman EtherType : 0x88a8
bvlan EtherType: 0x88b5
```

Displaying designated port for flooding control for rate-limited vMANs (BlackDiamond 12800 R-series switch only). The following is sample output from the `show vman <vlan_name>` command on the BlackDiamond 12800 R-series switch that is configured with vMAN ACLs and a designated port (the designated port information is on the ERL Designated Port line):

```
VMAN Interface with name vman110 created by user
  Admin State:      Enabled      Tagging:      802.1Q Tag 110
  Priority:         802.1P Priority 0
  Virtual router:   VR-Default
  IPv6:            None
  STPD:            None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disabled
  NetLogin:         Disabled
  ERL Designated Port: 6:2
  QosProfile:       QP1
  Ports:           4.           (Number of active ports=4)
    Untag:          *2:1,      *2:4
    Tag:            *6:1,      *6:2
  Flags:           (*) Active, (!) Disabled, (g) Load Sharing port
                  (b) Port blocked on the vlan, (a) Authenticated NetLogin port
                  (u) Unauthenticated NetLogin port, (m) Mac-Based port
```

Displaying the secondary ethertype and the secondary port list if present. The letter `g` in the port list indicates that the port is a LAG/Trunk port, the details of which can be seen using the `show port sharing` command.

```
(debug) BD-12804.3
# show vman ethertype
Vman Primary EtherType   : 0x9100
Vman Secondary EtherType : 0x8100
BVlan EtherType          : 0x88b5

Secondary EtherType ports : 6:2g 6:3
```

History

This command was first available in ExtremeXOS 11.0.

Information on IEEE 802.1ah was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

The 802.1ah information is available only on the BlackDiamond 10808 and 12800 series switch; the designated port information is available only on the BlackDiamond 12800 R-series switch.

unconfigure vman ethertype

```
unconfigure vman ethertype [ primary | secondary ]
```

Description

Changes the Ethernet type for the vMAN header back to the default value of 0x88A8 and type primary.

Syntax Description

primary	Specifies that the Ethertype is primary.
secondary	Specifies that the Ethertype is secondary.

Default

N/A.

Usage Guidelines

Use this command to restore the default values to the primary and secondary vMAN ethertypes. The primary vMAN ethertype defaults to 0x88A8 and the secondary defaults to 0x0. If you do not specify the primary or the secondary keyword, only the primary vMAN ethertype to return to its default value.



NOTE

Before unconfiguring the secondary vMAN ethertype, any secondary vMAN port must be changed to the primary vMAN ethertype; otherwise this command fails.

Example

The following command restores the primary vMAN ethertype to the default value:

```
unconfigure vman ethertype or unconfigure vman ethertype primary
```

The following command restores the secondary vMAN ethertype to the default value:

```
unconfigure vman ethertype secondary
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

The full command, with the addition of the `primary` and `secondary` keywords, is available on the BlackDiamond 12800 series switches only. The base command without the `primary` and `secondary` keywords, which applies only to the Primary MAN ethertype, is available on all platforms.

14 FDB Commands

This chapter describes commands for:

- Configuring FDB entries
- Displaying FDB entries

The switch maintains a database of all media access control (MAC) addresses received on all of its ports. It uses the information in this database to decide whether a frame should be forwarded or filtered.

Each FDB entry consists of the MAC address of the device, an identifier for the port on which it was received, and an identifier for the VLAN to which the device belongs. Frames destined for devices that are not in the FDB are flooded to all members of the VLAN.

The FDB has three types of entries:

- **Dynamic entries**—Initially, all entries in the database are dynamic. Entries in the database are removed (aged-out) if, after a period of time (aging time), the device has not transmitted. This prevents the database from becoming full of obsolete entries by ensuring that when a device is removed from the network, its entry is deleted from the database. Dynamic entries are deleted from the database if the switch is reset or a power off/on cycle occurs.
- **Static entries**—If the aging time is set to zero, all aging entries in the database are defined as static, nonaging entries. This means that they do not age, but they are still deleted if the switch is reset.
- **Permanent entries**—Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. The system administrator must create permanent entries. A permanent entry can either be a unicast or multicast MAC address. All entries entered through the command line interface (CLI) are stored as permanent.

Entries are added into the FDB in the following two ways:

- The switch can learn entries. The system updates its FDB with the source MAC address from a packet, the VLAN, and the port identifier on which the source packet is received.
- You can enter and update entries using the CLI.

clear fdb

```
clear fdb {<mac_addr> | ports <port_list> | vlan <vlan_name> | blackhole}
```

Description

Clears dynamic FDB entries that match the filter.

Syntax Description

mac_addr	Specifies a MAC address, using colon-separated bytes.
port_list	Specifies one or more ports or slots and ports.
vlan_name	Specifies a VLAN name.
blackhole	Specifies the blackhole entries.

Default

Clears all dynamic FDB entries.

Usage Guidelines

This command clears FDB entries based on the specified criteria. When no options are specified, the command clears all dynamic FDB entries.

Example

The following command clears any FDB entries associated with ports 4:3-4:5 on a modular switch:

```
clear fdb ports 4:3-4:5
```

The following command clears any FDB entries associated with VLAN *corporate*:

```
clear fdb vlan corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure fdb agingtime

```
configure fdb agingtime <seconds>
```

Description

Configures the FDB aging time for dynamic entries.

Syntax Description

seconds	Specifies the aging time in seconds. Range is 15 through 1,000,000. A value of 0 indicates that the entry should never be aged out.

Default

300 seconds.

Usage Guidelines

The range is 15 through 1,000,000 seconds.

If the aging time is set to zero, all aging entries in the database are defined as static, nonaging entries. This means that they do not age out, but non-permanent static entries can be deleted if the switch is reset.

The software flushes the FDB table once the aging timeout parameter is reached, even if the switch is running traffic and populating addresses in the FDB table.

Example

The following command sets the FDB aging time to 3,000 seconds:

```
configure fdb agingtime 3000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create fdbentry vlan ports

```
create fdbentry <mac_addr> vlan <vlan_name> ports <port_list>
```

Description

Creates a permanent static FDB entry, and optionally associates it with an ingress and/or egress QoS profile.

Syntax Description

mac_addr	Specifies a device MAC address, using colon-separated bytes.
vlan_name	Specifies a VLAN name associated with a MAC address.
port_list	Specifies one or more ports or slots and ports associated with the MAC address.

Default

N/A.

Usage Guidelines

Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. A permanent static entry can either be a unicast or multicast MAC address. After they have been created, permanent static entries stay the same as when they were created. If the same MAC address and VLAN is encountered on another virtual port that is not included in the permanent MAC entry, it is handled as a blackhole entry. The static entry is not updated when any of the following take place:

- A VLAN identifier (VLANid) is changed.
- A port is disabled.
- A port enters blocking state.
- A port goes down (link down).

A permanent static FDB entry is deleted when any of the following take place:

- A VLAN is deleted.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.

Permanent static entries are designated by “spm” in the flags field of the `show fdb` output. You can use the `show fdb` command to display permanent FDB entries.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS version 11.3, you can create FDB entries to multicast MAC addresses and list one or more ports. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations. The following restrictions apply:

- Not available for unicast MAC addresses.
- IGMP snooping rules take precedence over static multicast MAC addresses in the IP multicast range (01:00:5e:xx:xx:xx) unless IGMP snooping is disabled.

Example

The following command adds a permanent, static entry to the FDB for MAC address is 00 E0 2B 12 34 56, in VLAN marketing on slot 2, port 4 on a modular switch:

```
create fdbentry 00:E0:2B:12:34:56 vlan marketing port 2:4
```

The following example creates a multicast FDB entry, in VLAN black, on slot 1, ports 1, 2, and 4, on the BlackDiamond 8800 family of switches:

```
create fdbentry 01:00:00:00:00:01 vlan black port 1:1, 1:2, 1:4
```

History

This command was first available in ExtremeXOS 10.1.

The ability to create a multicast FDB with multiple entry ports was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

delete fdbentry

```
delete fdbentry [all | <mac_address> [vlan <vlan name>]
```

Description

Deletes one or all permanent FDB entries.

Syntax Description

all	Specifies all FDB entries.
mac_address	Specifies a device MAC address, using colon-separated bytes.
vlan	Specifies a VLAN.
vlan_name	Specifies the specific VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes a permanent entry from the FDB:

```
delete fdbentry 00:E0:2B:12:34:56 vlan marketing
```

The following example deletes all permanent entries from the FDB:

```
delete fdbentry all
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable flooding port

```
disable flooding [all_cast | broadcast | multicast | unicast] port
[<port_list> | all]
```

Description

Disables egress flooding on one or more ports. With the BlackDiamond 8800 series switch (formerly known as Aspen), SummitStack, and the Summit family of switches, you can further identify the type of packets to block flooding on the specified ports. With the BlackDiamond 10808 and 12800 series switches, you can block flooding only on all packets or no packets.

Syntax Description

all_cast	Specifies disabling egress flooding for <i>all</i> packets on specified ports.
broadcast	Specifies disabling egress flooding only for broadcast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
multicast	Specifies disabling egress flooding only for multicast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
unicast	Specifies disabling egress flooding only for unicast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines



NOTE

If an application requests specific packets on a specific port, those packets are not affected by the `disable flooding port` command.

Use this command to disable Layer 2 egress flooding on specified or all packets (only on all packets on the BlackDiamond 10808 and 12800 series switches). Disabling egress flooding aids the following:

- enhances security
- enhances privacy
- improves network performance

This is particularly useful when you are working on an edge device in the network. Limiting flooded egress packets to selected interfaces is also known as upstream forwarding.

**NOTE**

Disabling egress flooding with static MAC addresses can affect many protocols, such as IP and ARP among others.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. If that is the situation, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- FDB learning is independent of egress flooding. FDB learning and egress flooding can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded *to* that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded *to* that port.

BlackDiamond 8800 family of switches, SummitStack, and the Summit switch only. You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the BlackDiamond 8800 family of switches, SummitStack, and the Summit switch. The default behavior for the BlackDiamond 8800 family of switches, SummitStack, and the Summit is enabled egress flooding for all packet types.

BlackDiamond 10808 and 12800 series switches only. You must disable egress flooding on *all* packets on the specified port or ports. You cannot specify broadcast, unicast, or multicast packets; the egress flooding command applies to all packets. The default behavior for the BlackDiamond 10K and 12800 series switches is enabled egress flooding.

**NOTE**

When you disable egress flooding on the BlackDiamond 10808 and 12800 series switches, you also turn off broadcasting.

Example

The following command disables egress flooding on slot 4, ports 5 and 6 on a BlackDiamond 10808 switch:

```
disable flooding all_cast port 4:5-4:6
```

The following command disables unicast flooding on ports 10-12 on a Summit series switch:

```
disable flooding unicast port 10-27
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable flooding port

```
enable flooding [all_cast | broadcast | multicast | unicast] port
[<port_list> | all]
```

Description

Enables egress flooding on one or more ports. With the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches, you can further identify the type of packets to flood on the specified ports. With the BlackDiamond 10808 and 12800 series switches, you can flood only all packets or no packets.

Syntax Description

all_cast	Specifies enabling egress flooding for <i>all</i> packets on specified ports.
broadcast	Specifies enabling egress flooding only for broadcast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
multicast	Specifies enabling egress flooding only for multicast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
unicast	Specifies enabling egress flooding only for unicast packets. NOTE: This parameter is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines

Use this command to re-enable egress flooding that you previously disabled using the [disable flooding port](#) command.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. If that is the situation, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- FDB learning is independent of egress flooding. FDB learning and egress flooding can be enabled or disabled independently.

- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded *to* that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded *to* that port.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches. The default behavior for the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches is enabled egress flooding for all packet types.

BlackDiamond 10808 and 12800 series switches only. You must disable egress flooding on *all* packets on the specified port or ports. You cannot specify broadcast, unicast, or multicast packets; the egress flooding command applies to all packets. The default behavior for the BlackDiamond 10808 and 12800 series switches is enabled egress flooding.



NOTE

When you disable egress flooding on the BlackDiamond 10808 and 12800 series switches, you also turn off broadcasting.

Example

The following command enables egress flooding on slot 1, ports 1 and 2 on a BlackDiamond 10808 switch:

```
enable flooding all_cast port 1:1-1:2
```

The following command enables unicast flooding on ports 13-17 on a Summit series switch:

```
enable flooding unicast port 13-17
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show fdb

```
show fdb {<mac_addr> {netlogin [all | mac-based-vlans]}| permanent
{netlogin [all | mac-based-vlans]} | ports <port_list> {netlogin [all |
mac-based-vlans]}| vlan <vlan_name> {netlogin [all | mac-based-vlans]} |
stats | netlogin {all | mac-based-vlans}}}
```

Description

Displays FDB entries.

Syntax Description

mac_addr	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
netlogin	Displays FDBs created as a result of the netlogin process.
all	Displays all FDBs created as a result of the netlogin process.
mac-based-vlans	Displays all netlogin MAC-based VLAN FDB entries. NOTE: This parameter is supported only for the Summit family of switches, SummitStack, and the BlackDiamond 8800 series switches. See Chapter 14, “Network Login Commands,” for more information on netlogin.
permanent	Displays all permanent entries, including the ingress and egress QoS profiles.
port_list	Displays the entries for one or more ports or ports and slots.
vlan_name	Displays the entries for a specific VLAN.
stats	Displays the number of static, permanent, dynamic, and dropped FDB entries, as well as the FDB aging time.

Default

All.

Usage Guidelines

Displays FDB entries as specified, or displays all FDB entries.

The show output displays the following information:

Mac	The MAC address that defines the entry.
Vlan	The VLAN for the entry.
Age	The age of the entry, in seconds (does not appear if the keyword permanent is specified). The age parameter does not display for the backup MSM on modular switches.
Use	The number of IP FDB entries that use this MAC address as a next hop or last hop (does not appear if the keyword permanent is specified).

Flags	Flags that define the type of entry: • B - Egress Blackhole • b - Ingress Blackhole • d - Dynamic • s - Static • p - Permanent • m - MAC • n - NetLogin • l - lockdown MAC • M - Mirror • i - an entry also exists in the IP FDB • x - an entry also exists in the IPX FDBs • v - NetLogin MAC-Based VLAN (only supported on the Summit switch, SummitStack, and the BlackDiamond 8800 family of switches)
Port List	The ports on which the MAC address has been learned.

Example

The following command displays information about all the entries in the FDB:

```
show fdb
```

It produces output similar to the following:

```

Mac                Vlan      Age  Flags      Port / Virtual Port List
-----
00:50:04:d6:e4:92   serv(4092) 0009 d m       3:5
00:80:c8:ca:e4:d9   gtag101(0101) 0000 d mi    3:4
00:80:c8:ca:e4:d9   serv(4092) 0012 d m     gv:1.1.1.46

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
        x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
Blackhole,
        b - Ingress Blackhole, v - MAC-Based VLAN.

Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
FDB Aging time: 300
FDB VPLS Aging time: 300

```

History

This command was first available in ExtremeXOS 10.1.

The stats and netlogin parameters were first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

This chapter describes commands for:

- Creating and deleting virtual routers
- Configuring and managing virtual routers
- Displaying information about virtual routers

Virtual Routers

ExtremeXOS supports virtual routers. This capability allows a single physical switch to be split into multiple virtual routers. This feature separates the traffic forwarded by a virtual router from the traffic on a different virtual router.

There are two types of virtual routers in an ExtremeXOS system:

- System virtual routers
These are the special virtual routers created by ExtremeXOS during system boot up, and they cannot be deleted or renamed. There are a total of three of these special virtual routers in the ExtremeXOS system.
- User virtual routers
These are the virtual routers created and named by users.



NOTE

User virtual routers are supported only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

Each virtual router maintains a separate logical forwarding table, which allows the virtual routers to have overlapping address spaces. Because each virtual router maintains its own separate routing information, packets arriving on one virtual router will never be switched to another.

With multiple virtual routers contained on a single physical switch, some commands in ExtremeXOS now require you to specify to which virtual router the command applies. For example, when you use the ping command, you must specify from which virtual router the ping packets are generated. Many commands that deal with switch management use the management virtual router by default.



NOTE

The term virtual router is also used with the Virtual Router Redundancy Protocol (VRRP). VRRP uses the term to refer to a single virtual router that spans more than one physical router, which allows multiple switches to provide redundant routing services to users.

create virtual-router

```
create virtual-router <vr-name>
```

Description

Creates a user virtual router.

Syntax Description

vr-name	Specifies the name of the user virtual router.

Default

N/A.

Usage Guidelines

This command creates a new user virtual router. The three default system virtual routers, VR-Mgmt, VR-Control, and VR-Default always exist and cannot be deleted or renamed.

A virtual router is identified by a name (up to 32 characters long). The name must be unique among the VLAN and virtual router names on the switch. For backward compatibility, you cannot name a virtual router VR-0, VR-1, or VR-2, as they were the original names of the system virtual routers. Virtual router names are case insensitive.

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.-

Example

The following command creates the virtual router *vr-acme*:

```
create virtual-router vr-acme
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

configure vr add protocol

```
configure vr <vr-name> add protocol <protocol-name>
```

Description

Starts the Layer 3 protocol on a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.
protocol-name	Specifies the layer 3 protocol.

Default

N/A.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to start the layer 3 protocol specified on the virtual router. The choices for `protocol-name` are:

- RIP
- OSPF
- BGP
- PIM

You cannot add or delete protocols from the system virtual routers, VR-Mgmt, VR-Control, VR-Default.

Example

The following command starts RIP on the virtual router *vr-acme*:

```
configure vr vr-acme add protocol rip
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

configure vr add ports

```
configure vr <vr-name> add ports <portlist>
```

Description

Assigns a list of ports to the virtual router specified.

Syntax Description

vr-name	Specifies the name of the virtual router.
portlist	Specifies the ports to add to the virtual router.

Default

By default, all ports are assigned to the virtual router, VR-Default.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to assign ports to a virtual router. Since all ports are initially assigned to VR-Default, you may need to delete the desired ports first from the virtual router where they reside, before you add them to the desired virtual router.

Example

The following command adds all the ports on slot 2 to the virtual router *vr-acme*:

```
configure vr vr-acme add ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

configure vr delete protocol

```
configure vr <vr-name> delete protocol <protocol-name>
```

Description

Stops and removes the layer 3 protocol on a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.
protocol-name	Specifies the layer 3 protocol.

Default

N/A.

Usage Guidelines

You cannot add or delete protocols from the system virtual routers, VR-Mgmt, VR-Control, VR-Default.

Example

The following command shutsdowns and removes RIP from the virtual router *vr-acme*:

```
configure vr vr-acme delete protocol rip
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

configure vr delete ports

```
configure vr <vr-name> delete ports <portlist>
```

Description

Removes a list of ports from the virtual router specified.

Syntax Description

vr-name	Specifies the name of the virtual router.
portlist	Specifies the ports to remove from the virtual router.

Default

By default, all ports are assigned to the virtual router, VR-Default.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to remove ports from a virtual router. Since all ports are initially assigned to VR-Default, you may need to delete the desired ports first from the virtual router where they reside, before you add them to the desired virtual router.

Example

The following command removes all the ports on slot 2 from the virtual router *vr-acme*:

```
configure vr vr-acme delete ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

delete virtual-router

```
delete virtual-router <vr-name>
```

Description

Deletes a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.

Default

N/A.

Usage Guidelines

Only user virtual routers can be deleted. When a virtual router gets deleted, all of the VLANs in the virtual router will be deleted. All of the ports assigned to this virtual router will be deleted and made available to assign to other virtual routers. Any routing protocol that is running on the virtual router will be shutdown and deleted gracefully.

Example

The following command creates the virtual router *vr-acme*:

```
delete virtual-router vr-acme
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

show virtual-router

```
show virtual-router {<vr-name>}
```

Description

Displays information about the virtual routers.

Syntax Description

vr-name	Specifies the name of the virtual router.

Default

N/A.

Usage Guidelines

During system boot up, ExtremeXOS creates three system virtual routers: VR-Mgmt, VR-Control, and VR-Default (previous to release 11.0 these virtual routers were named VR-0, VR-1, and VR-2, respectively). The following defines each system virtual router:

- The management port on both the primary and backup MSMs and the VLAN *mgmt* belong to VR-Mgmt.
- Internal system operations use VR-Control.
- The default VLAN belongs to VR-Default.

Beginning with release 11.0, you can create additional virtual routers, called user virtual routers. User virtual routers are created without any routing protocols, so the protocols must be added. The protocols on the system virtual routers are predefined and cannot be changed.

The output displays, in tabular format, the:

- Name of the virtual router
- Number of VLANs that belong to that virtual router
- Number of ports that belong to that virtual router
- Which routing protocols have been added to that virtual router

When you specify a particular virtual router, the output displays:

- The number of ports
- A list of ports
- The protocols configured
- The name of the process supporting the protocol on that virtual router

Example

The following command displays the virtual router configurations on the switch:

```
show virtual-router
```

The following is sample output from this command:

```
-----
Virtual Router      Number of   Number of   Flags
                   Vlans       Ports
-----
VR-Control          0           0           ----
VR-Default          4           776         bopr
VR-Mgmt             1           0           ----
-----
Flags : Routing protocols configured on the virtual router
        (b) BGP, (o) OSPF, (r) RIP, (p) PIM
```

The following command displays the virtual router *helix*:

```
show virtual-router helix
```

The following is sample output from this command:

```
Virtual router      : helix
No of ports         : 2
Port List           : 3:1-2
Protocols Configured:
    Protocol: BGP,   Process Name: bgp-4
    Protocol: RIP,   Process Name: rip-4
```

History

A command similar to this command was available in ExtremeXOS 10.1 (`show vr`).

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

virtual-router

```
virtual-router {<vr-name>}
```

Description

Changes the virtual router domain context.

Syntax Description

vr-name	Specifies the name of the virtual router.

Default

N/A.

Usage Guidelines

Use this command to change the virtual router context for subsequent commands. When you issue the command, the prompt changes to reflect the virtual router domain. Configuration commands for layer 3 routing protocols, creating VLANs, and deleting VLANs apply only to the current virtual router domain.

Under a virtual router configuration domain, any virtual router commands are applied only to that virtual router. The virtual router commands consist of all the BGP, OSPF, PIM and RIP commands, and the commands listed in [Table](#) .

Table 20: Virtual router commands

[enable disable] ipforwarding
clear iparp *
clear counters iparp *
configure iparp *
configure iparp [add delete] *
[enable disable] iparp *
show iparp *
configure iproute [add delete] *
show iproute *
show ipstats *
rtlookup
create [vlan vman] <vlan-name>
[enable disable] igmp
[enable disable] igmp snooping *
[enable disable] ipmcforwarding
show igmp

Table 20: Virtual router commands (Continued)

show igmp snooping
show igmp group
show igmp snooping cache

* means that other commands are available with these listed.

The virtual router configuration domain simplifies configuration because you do not have to specify the virtual router for each individual protocol configuration command. The current configuration domain is indicated in the command line interface (CLI) prompt.

For example, if you wish to configure OSPF for the user virtual router *vr-manufacturing*, you would change the virtual router context to that of *vr-manufacturing*. All the subsequent OSPF commands would apply to that virtual router, unless the context is changed again.

A virtual router is identified by a name (up to 32 characters long). The name must be unique among the VLAN and virtual router names on the switch. For backward compatibility, you cannot name a virtual router VR-0, VR-1, or VR-2. Virtual router names are case insensitive.

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

Example

The following command changes the virtual router context to *vr-acme*:

```
virtual-router vr-acme
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

16 Policy Manager Commands

This chapter describes commands for:

- Creating and configuring policy files for IP access lists (ACLs)
- Creating and configuring policy files for routing policies

Policies are a generalized category of features that impact forwarding and route forwarding decisions. Access policies are used primarily for security and quality of service (QoS) purposes.

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, for the BlackDiamond 8800 family and Summit series switches only, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses layer 3 router boundaries, but it is possible to use access lists within a layer 2 VLAN. Extreme products are capable of performing this function with no additional configuration.

Routing policies are used to control the advertisement or recognition of routes from routing protocols, such as RIP, OSPF, or BGP. Routing policies can be used to 'hide' entire networks or to trust only specific sources for routes or ranges of routes. The capabilities of routing policies are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.



NOTE

Although ExtremeXOS does not prohibit mixing ACL and routing type entries in a policy file, it is strongly recommended that you do not mix the entries, and you use separate policy files for ACL and routing policies.

check policy

```
check policy <policy-name> {access-list}
```

Description

Checks the syntax of the specified policy.

Syntax Description

policy-name	Specifies the policy to check.
access-list	Specifies that an access list specific check is performed.

Default

N/A.

Usage Guidelines

Use this command to check the policy syntax before applying it. If any errors are found, the line number and a description of the syntax error are displayed. A policy that contains syntax errors will not be applied.

This command can only determine if the syntax of the policy file is correct and can be loaded into the policy manager database. Since a policy can be used by multiple applications, a particular application may have additional constraints on allowable policies.

Example

The following example checks the syntax of the policy *zone5*:

```
check policy zone5
```

If no syntax errors are discovered, the following message is displayed:

```
Policy file check successful.
```

History

This command was available in ExtremeXOS 10.1.

The success message and the access-list keyword was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

check policy attribute

```
check policy attribute {<attr>}
```

Description

Displays the syntax of the specified policy attribute.

Syntax Description

attr	Specifies the attribute check.

Default

N/A.

Usage Guidelines

Use this command to display the syntax of policy attributes. The command displays any additional keywords to use with this attribute, and the types of values expected.

Policy attributes are used in the rule entries that make up a policy file.

For each attribute, this command displays which applications use the attribute, and whether the attribute is a match condition or a set (action, action modifier) condition.

The current applications are:

- ACL—access-lists
- RT—routing profiles, route maps
- CLF—CLEAR-Flow

The syntax display does not show the text synonyms for numeric entries. For example, the `icmp-type` match condition allows you to specify either an integer or a text synonym for the condition. Specifying `icmp-type 8` or `icmp-type echo-request` are equivalent, but the syntax display shows only the numeric option.



NOTE

The syntax displayed is used by the policy manager to verify the syntax of policy files. The individual applications are responsible for implementing the individual attributes. Inclusion of a particular policy attribute in this command output does not imply that the attribute has been implemented by the application. See the documentation of the particular application for detailed lists of supported attributes.

Example

The following example displays the syntax of the policy attribute `icmp-type`:

```
check policy attribute icmp-type
```

The following is sample output for this command:

```
( match ) ( ACL )  
icmp-type <uint32 val>
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

edit policy

```
edit policy <filename>
```

Description

Edits a policy text file.

Syntax Description

filename	Specifies the filename of the policy text file.

Default

N/A.

Usage Guidelines

This command edits policy text files that are on the switch. All policy files use “.pol” as the filename extension, so to edit the text file for the policy *boundary* use *boundary.pol* as the filename. If you specify the name of a file that does not exist, you will be informed and the file will be created.

This command spawns a VI-like editor to edit the named file. For information on using VI, if you are not familiar with it, do a web search for “VI editor basic information”, and you should find many resources. The following is only a short introduction to the editor.

Edit operates in one of two modes; command and input. When a file first opens, you are in the command mode. To write in the file, use the keyboard arrow keys to position your cursor within the file, then press one of the following keys to enter input mode:

- i - To insert text ahead of the initial cursor position
- a- To append text after the initial cursor position

To escape the input mode and return to the command mode, press the Escape key.

There are several commands that can be used from the command mode. The following are the most commonly used:

- dd - To delete the current line
- yy - To copy the current line
- p - To paste the line copied
- :w - To write (save) the file
- :q - To quit the file if no changes were made
- :q! - To forcefully quit the file without saving changes
- :wq - To write and quit the file

Refresh Policy. After you have edited the text file for a policy that is currently active, you will need to refresh the policy if you want the changes to be reflected in the policy database. When you refresh the

policy, the text file is read, the syntax is checked, the policy information is added to the policy manager database, and the policy then takes effect. Use the following command to refresh a policy:

```
refresh policy <policy-name>
```

If you just want to check to be sure the policy contains no syntax errors, use the following command:

```
check policy <policy-name> {access-list}
```

Example

The following command allows you to begin editing the text file for the policy *boundary*:

```
edit policy boundary.pol
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

refresh policy

```
refresh policy <policy-name>
```

Description

Refreshes the specified policy.

Syntax Description

policy-name	Specifies the policy to refresh.

Default

N/A.

Usage Guidelines

Use this command when a new policy file for a currently active policy has been downloaded to the switch, or when the policy file for an active policy has been edited. This command reprocesses the text file and updates the policy database.

In releases previous to ExtremeXOS 11.4, when ACLs were refreshed, all the ACL entries were removed, and new ACL entries were created to implement the newly applied policy. Beginning in release 11.4, the policy manager uses Smart Refresh to update the ACLs. When a change is detected, only the ACL changes needed to modify the ACLs are sent to the hardware, and the unchanged entries remain. This behavior avoids having to blackhole packets because the ACLs have been momentarily cleared. Smart Refresh works well for minor changes, however, if the changes are too great, the refresh reverts to the earlier behavior. To take advantage of Smart Refresh, disable access-list refresh blackholing by using the command:

```
disable access-list refresh blackhole
```

If you attempt to refresh a policy that cannot take advantage of Smart Refresh, you will receive a message similar to the following if blackholing is enabled:

```
Incremental refresh is not possible given the configuration of policy <name>. Note,
the current setting for Access-list Refresh Blackhole is Enabled.
Would you like to perform a full refresh? (Yes/No) [No]:
```

and if blackholing is not enabled:

```
Incremental refresh is not possible given the configuration of policy <name>. Note,
the current setting for Access-list Refresh Blackhole is Disabled.
WARNING: If a full refresh is performed, it is possible packets that should be denied
may be forwarded through the switch during the time the access list is being
installed.
Would you like to perform a full refresh? (Yes/No) [No]:
```

If you attempt to refresh a policy that is not currently active, you will receive an error message.

For an ACL policy, the command is rejected if there is a configuration error or hardware resources are not available.

Example

The following example refreshes the policy *zone5*:

```
refresh policy zone5
```

History

This command was first available in ExtremeXOS 11.0.

Smart Refresh was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show policy

```
show policy {<policy-name> | detail}
```

Description

Displays the specified policy.

Syntax Description

policy-name	Specifies the policy to display.
detail	Show the policy in detail.

Default

If no policy name is specified, all policies are shown

Usage Guidelines

Use this command to display which clients are using the specified policy. The detail option displays the rules that make up the policy.

Example

The following example displays the policy *zone5*:

```
show policy zone5
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

17 ACL Commands

This chapter describes commands for:

- Creating and configuring IP access lists (ACLs)

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, for the BlackDiamond 8800 series, SummitStack, and Summit family switches only, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses layer 3 router boundaries, but it is possible to use access lists within a layer 2 VLAN. Extreme products are capable of performing this function with no additional configuration.



NOTE

Although ExtremeXOS does not prohibit mixing ACL and routing type entries in a policy file, it is strongly recommended that you do not mix the entries, and you use separate policy files for ACL and routing policies.

clear access-list counter

```
clear access-list {dynamic} counter {<countername>} {any | ports <portlist>
| vlan <vlanname>} {ingress | egress}
```

Description

Clears the specified access list counters.

Syntax Description

dynamic	Specifies that the counter is from a dynamic ACL.
countername	Specifies the ACL counter to clear.
any	Specifies the wildcard ACL.
portlist	Specifies to clear the counters on these ports.
vlanname	Specifies to clear the counters on the VLAN.
ingress	Clear the ACL counter for packets entering the switch on this interface.
egress	Clear the ACL counter for packets leaving the switch from this interface (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default direction is ingress; the default ACL type is non-dynamic.

Usage Guidelines

Use this command to clear the ACL counters. If you do not specify an interface, or the *any* option, you will clear all the counters.

Example

The following example clears all the counters of the ACL on port 2:1:

```
clear access-list counter port 2:1
```

The following example clears the counter *counter2* of the ACL on port 2:1

```
clear access-list counter counter2 port 2:1
```

History

This command was available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress and dynamic options were first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

clear access-list meter

```
clear access-list meter {<metername>} [any | ports <portlist> | vlan <vlanname>]
```

Description

Clears the specified access list meters.

Syntax Description

metername	Specifies the ACL meter to clear.
portlist	Specifies to clear the counters on these ports.
vlanname	Specifies to clear the counters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to clear the out-of-profile counters associated with the meter configuration.

Example

The following example clears all the out-of-profile counters for the meters of the ACL on port 2:1:

```
clear access-list meter port 2:1
```

The following example clears the out-of-profile counters for the meter *meter2* of the ACL on port 2:1

```
clear access-list meter meter2 port 2:1
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 family, SummitStack, and Summit switches.

configure access-list

```
configure access-list <aclname> [any | ports <portlist> | vlan <vlanname>]
{ingress | egress}
```

Description

Configures an access list to the specified interface.

Syntax Description

aclname	Specifies the ACL name. The name can be from 1-32 characters long.
any	Specifies that this ACL is applied to all interfaces as the lowest precedence ACL.
portlist	Specifies the ports on which this ACL is applied.
vlanname	Specifies the VLAN on which this ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default direction is ingress.

Usage Guidelines

The access list applied in this command is contained in a text file created either externally to the switch or using the [edit policy](#) command. The file is transferred to the switch using TFTP before it is applied to the ports. The ACL name is the file name without its “.pol” extension. For example, the ACL *blocknetfour* would be in the file *blocknetfour.pol*. For more information on policy files, see the *ExtremeXOS Concepts Guide*.

Specifying the keyword *any* applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to it, and is also applied to packets that do not match the ACL applied to the interface.

Example

The following command configures the ACL *test* to port 1:2 at ingress:

```
configure access-list test ports 1:2
```

The following command configures the ACL *mydefault* as the wildcard ACL:

```
configure access-list mydefault any
```

The following command configures the ACL *border* as the wildcard egress ACL:

```
configure access-list border any egress
```

History

This command was available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress options was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

configure access-list add

```
configure access-list add <dynamic_rule> [ [[first | last]{priority
<p_number>}] | [[before | after] <rule>] | [priority <p_number>]] [any | vlan
<vlanname> | ports <portlist>] {ingress | egress} {zone <zone>}
{application <appl_name>}
```

Description

Configures a dynamic ACL rule to the specified interface, sets the priority and zone for the ACL, and binds it to an application.

Syntax Description

dynamic_rule	Specifies a dynamic ACL rule.
first	Add the new dynamic rule before all existing dynamic rules in the area assigned to the application.
last	Add the new dynamic rule after all existing dynamic rules in the area assigned to the application.
priority <p_number>	Specifies the priority number of the rule. The range is from 0 (highest priority) to 7 (lowest priority) within the area assigned to the application.
before <rule>	Add the new dynamic rule before an existing dynamic rule.
after <rule>	Add the new dynamic rule after an existing dynamic rule.
any	Specifies that this ACL is applied to all interfaces.
vlanname	Specifies the VLAN on which this ACL is applied.
portlist	Specifies the ports on which this ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface (BlackDiamond 10808 and BlackDiamond 12800 series only).
zone	Specifies the ACL zone for the rule.
application <appl_name>	Specifies an application to which the ACL is bound.

Default

The default direction is ingress.

Usage Guidelines

The dynamic rule must first be created before it can be applied to an interface. Use the following command to create a dynamic rule:

```
create access-list <dynamic-rule> <conditions> <actions> {non-permanent} {application
<appl_name>}
```

When a dynamic ACL rule is applied to an interface, you will specify its precedence among any previously applied dynamic ACLs. All dynamic ACLs have a higher precedence than any ACLs applied through ACL policy files.

Specifying the keyword `any` applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to them, and is also applied to packets that do not match the ACL applied to the interface.

The `priority` keyword can be used to set the priority of a rule within a zone, or to create a sub-zone within an application's zone. For example, to place ACLs into three sub-zones within the CLI application, you can use three priority numbers, such as 2, 4, and 7.

Configuring priority number 1 is the same as configuring *first* priority. Configuring priority number 8 is the same as configuring *last* priority.

Example

The following command applies the dynamic ACL *icmp-echo* as the first (highest precedence) dynamic ACL to port 1:2 at ingress:

```
configure access-list add icmp-echo first ports 1:2
```

The following command applies the dynamic ACL *udpacl* to port 1:2, with a higher precedence than rule *icmp-echo*:

```
configure access-list add udpacl before icmp-echo ports 1:2
```

History

This command was first available in ExtremeXOS 11.3.

The application option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

configure access-list delete

```
configure access-list delete <dynamic_rule> [any | vlan <vlanname> | ports
<portlist> | all] {ingress | egress} {application <appl_name>}
```

Description

Removes a dynamic ACL rule from the specified interface or application.

Syntax Description

<dynamic_rule>	Specifies a dynamic ACL rule.
any	Delete this ACL as the wildcard ACL.
vlanname	Specifies the VLAN on which this ACL is deleted.
portlist	Specifies the ports on which this ACL is deleted.
all	Delete this ACL from all interfaces.
ingress	Delete the ACL for packets entering the switch on this interface.
egress	Delete the ACL for packets leaving the switch from this interface (BlackDiamond 10808 and BlackDiamond 12800 series only).
application <appl_name>	Deletes this ACL from the specified application.

Default

The default direction is ingress.

Usage Guidelines

Specifying the keyword `all` removes the ACL from all interfaces it is used on.

Use the `application` keyword to remove the ACL from a specific application.

Example

The following command removes the dynamic ACL *icmp-echo* from the port 1:2:

```
configure access-list delete icmp-echo ports 1:2
```

History

This command was first available in ExtremeXOS 11.3.

The application option was added in ExtremeXOS 11.6

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

configure access-list zone

```
configure access-list zone <name> zone-priority <number>

configure access-list zone <name> move-application <appl-name> to-zone
<name> application-priority <number>

configure access-list zone <name> {add} application <appl-name>
application_priority <number>

configure access-list zone <name> delete application <appl-name>
```

Description

Configures the priority of a zone; moves an application from one zone to another at a specified priority; adds an application to a zone with a specified priority, or changes the priority of an application within a zone; deletes an application from a zone.

Syntax Description

<name>	Specifies a a zone name.
zone-priority <number>	Sets the priority of the zone.
move-application <appl-name>	Specifies the name of an application to be moved.
to-zone <name>	Specifies the zone to which the application is moved.
application-priority <number>	Sets the priority of the application within the zone. The range is from 0 (highest priority) to 7 (lowest priority).
add	Adds an application to a zone at a specified priority.
application <appl_name>	Specifies the application to be added to the zone.
application_priority <number>	Sets the priority of a new or existing application within a zone. The range is from 0 (highest priority) to 7 (lowest priority).

Default

N/A.

Usage Guidelines

To configure the priority of a specific zone, use the syntax:

```
configure access-list zone <name> zone-priority <number>
```

To move an application from one zone to another, and set its priority in the new zone, use the syntax:

```
configure access-list zone <name> move-application <appl-name> to-zone
<name> application-priority <number>
```

To add an application to a zone and specify its priority or to change the priority of an application within a zone, use the syntax:

```
configure access-list zone <name> {add} application <appl-name>  
application_priority <number>
```

To delete an application from a zone, use the syntax:

```
configure access-list zone <name> delete application <appl-name>
```

Example

The following command adds the CLI application to the zone *myzone* at a priority of 6:

```
configure access-list zone myzone add cli application-priority 6
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

create access-list

```
create access-list <dynamic-rule> <conditions> <actions> {non-permanent}
{application <appl_name>}
```

Description

Creates a dynamic ACL

Syntax Description

dynamic-rule	Specifies the dynamic ACL name. The name can be from 1-32 characters long.
conditions	Specifies the match conditions for the dynamic ACL.
actions	Specifies the actions for the dynamic ACLs.
non-permanent	Specifies that the ACL is not permanent.
application <appl_name>	Specifies an application to which the ACL belongs.

Default

By default, ACLs are permanent, and belong to the CLI application.

Usage Guidelines

This command creates a dynamic ACL rule. Use the `configure access-list add` command to apply the ACL to an interface.

The `conditions` parameter is a quoted string of match conditions, and the `actions` parameter is a quoted string of actions. Multiple match conditions or actions are separated by semi-colons. A complete listing of the match conditions and actions is in the ExtremeXOS Concepts Guide, in [Chapter 12, "Access Lists \(ACLs\)."](#)

Dynamic ACL rule names must be unique, but can be the same as used in a policy-file based ACL. Any dynamic rule counter names must be unique.

By default, ACL rules are saved when the `save` command is executed, and persist across system reboots. Configuring the optional keyword `non-permanent` means the ACL will not be saved.

You can configure an ACL to belong to a specific application, using the optional `application <applName>` keyword and argument.

Example

The following command creates a dynamic ACL that drops all ICMP echo-request packets on the interface:

```
create access-list icmp-echo "protocol icmp;icmp-type echo-request" "deny"
```

The created dynamic ACL will take effect after it has been configured on the interface. The previous example creates a dynamic ACL named *icmp-echo* that is equivalent to the following ACL policy file entry:

```
entry icmp-echo {
    if {
        protocol icmp;
        icmp-type echo-request;
    } then {
        deny;
    }
}
```

The following command creates a dynamic ACL that accepts all the UDP packets from the 10.203.134.0/24 subnet that are destined for the host 140.158.18.16, with source port 190 and a destination port in the range of 1200 to 1250:

```
create access-list udpacl "source-address 10.203.134.0/24;destination-address
140.158.18.16/32;protocol udp;source-port 190;destination-port 1200 - 1250;"
"permit"
```

The previous example creates a dynamic ACL entry named *udpac1* that is equivalent to the following ACL policy file entry:

```
entry udpac1 {
    if {
        source-address 10.203.134.0/24;
        destination-address 140.158.18.16/32;
        protocol udp;
        source-port 190;
        destination-port 1200 - 1250;
    } then {
        permit;
    }
}
```

History

This command was first available in ExtremeXOS 11.3.

The non-permanent and application options were added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

create access-list zone

```
create access-list zone <name> zone-priority <number>
```

Description

Creates a dynamic ACL zone, and sets the priority of the zone.

Syntax Description

<name>	Specifies the dynamic ACL zone name. The name can be from 1-32 characters long.
zone-priority <number>	Specifies priority of the zone. The range is from 1 (highest priority) to 4294967295 (lowest priority).

Default

The denial of service, system, and security zones are configured by default, and cannot be deleted.

Usage Guidelines

This command creates a dynamic ACL zone. You can configure the priority of the zone in relation to the default zones or to other configured zones.

Example

The following command creates a new zone, called myzone, with a priority of 2:

```
create access-list myzone zone-priority 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete access-list

```
delete access-list <dynamic-rule>
```

Description

Deletes a dynamic ACL

Syntax Description

dynamic-rule	Specifies the dynamic ACL name.

Default

N/A.

Usage Guidelines

This command deletes a dynamic ACL rule. Before you delete a dynamic ACL, it must be removed from any interfaces it is applied to. Use the `configure access-list delete` command to remove the ACL from an interface.

Example

The following command deletes the dynamic ACL *icmp-echo*:

```
delete access-list icmp-echo
```

History

This command was first available in ExtremeXOS 11.3.

The `non-permanent` and `application` options were added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete access-list zone

```
delete access-list zone <name>
```

Description

Deletes an ACL zone.

Syntax Description

<name>	Specifies the zone name.

Default

N/A.

Usage Guidelines

This command deletes an ACL zone. You must remove all applications from a zone before you can delete the zone. To delete an application from a zone, use the command `configure access-list zone <name> delete application <appl-name>`

You cannot delete the default zones.

Example

The following command deletes the zone *my_zone*:

```
delete access-list zone my_zone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable access-list permit to-cpu

```
disable access-list permit to-cpu
```

Description

Allows special packets to be blocked by low priority ACLs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows ACLs to deny certain special packets from reaching the CPU, even if the packets match ACLs that would otherwise deny them. The special packets include STP and EAPS BPDUs, and ARP replies for the switch.

When this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, the packets will still be denied if there is a higher precedence entry that permits the packets.

To enable this feature, use the following command:

```
enable access-list permit to-cpu
```

Example

The following command enables ACLs to deny STP BPDU packets from reaching the switch CPU:

```
disable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available only on the BlackDiamond 8800 original series modules and Summit X450 series switches.

disable access-list refresh blackhole

```
disable access-list refresh blackhole
```

Description

Disables blackholing of packets during ACL refresh.

Syntax Description

This command has no arguments or variables.

Default

The feature is enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this feature provides that any packets arriving during the refresh will be blackholed.

If you disable this feature, the ACLs will be refreshed as described in the [refresh policy](#) command.

To enable this feature, use the following command:

```
enable access-list refresh blackhole
```

Example

The following command disables dropping of packets during an ACL refresh:

```
disable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable access-list permit to-cpu

```
enable access-list permit to-cpu
```

Description

Enables control packets to reach CPU, even if an ACL would deny them.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows control packets to reach the CPU, even if the packets match ACLs that would otherwise deny them. The control packets include STP and EAPS BPDUs, and ARP replies for the switch.

If this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, when this feature is disabled, the packets will still be denied if there is a higher precedence entry that permits the packets.

To disable this feature, use the following command:

```
enable access-list permit to-cpu
```

Example

The following command enables STP BPDU packets to reach the switch CPU, despite any ACL:

```
enable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available only on the BlackDiamond 8800 family, SummitStack, and Summit switches.

enable access-list refresh blackhole

```
enable access-list refresh blackhole
```

Description

Enables blackholing of packets during ACL refresh.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this command provides that any packets arriving during the refresh will be blackholed. As the ACL is being refreshed, packets may arrive while the ACL is in an indeterminate state, and packets may be permitted that otherwise are dropped. This feature protects the switch during an ACL refresh. To disable this feature, use the following command:

```
disable access-list refresh blackhole
```

Example

The following command enables dropping of packets during an ACL refresh:

```
enable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show access-list

```
show access-list {any | ports <portlist> | vlan <vlanname>} {ingress | egress}
```

Description

Displays the ACLs configured on an interface.

Syntax Description

aclname	Specifies the ACL name. The name can be from 1-32 characters long.
any	Specifies the wildcard ACL.
portlist	Specifies which ports' ACLs to display.
vlanname	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default is to display all interfaces, ingress.

Usage Guidelines

The ACL with the port and VLAN displayed as an asterisk (*) is the wildcard ACL.

If you do not specify an interface, the policy names for all the interfaces are displayed, except that dynamic ACL rule names are not displayed. To display dynamic ACLs use the following commands:

```
show access-list dynamic
show access-list dynamic rule <rule> {detail}
```

If you specify an interface, all the policy entries, and dynamic policy entries are displayed.

Example

The following command displays all the interfaces configured with an ACL:

```
show access-list
```

The output from this command is similar to:

Vlan Name	Port	Policy Name	Dir	Rules	Dyn Rules
=====					
*	3:6	TCP_flag	ingress	3	2
*	3:8	qos_hongkong	ingress	3	0
*	2:1	tc_2.4	ingress	4	0
*	2:7	tcp	ingress	1	0
v1	*	tcp	ingress	1	0
*	*	firewall1	ingress	2	1

The following command displays the ingress access list entries configured on the VLAN v1006:

```
show access-list v1006 ingress
```

The output from this command is similar to the following:

```
# RuleNo 1
entry dacl13 {          #Dynamic Entry
if match all {
    ethernet-destination-address 00:01:05:00:00:00 ;
} then {
    count c13 ;
    redirect 1.1.5.100 ;
} }

# RuleNo 2
entry dacl14 {          #Dynamic Entry
if match all {
    ethernet-source-address 00:01:05:00:00:00 ;
} then {
    count c14 ;
    qosprofile qp7 ;
} }

# RuleNo 3
entry dacl13 {
if match all {
    ethernet-destination-address 00:01:05:00:00:00 ;
} then {
    count c13 ;
    redirect 1.1.5.100 ;
} }
```

History

This command was available in ExtremeXOS 10.1.

The <aclname> option was removed in ExtremeXOS 11.1.

The ingress, egress, any, ports, and vlan options were added in ExtremeXOS 11.3

Platform Availability

This command is available on all platforms.

show access-list configuration

```
show access-list configuration
```

Description

Displays the ACL configuration. Currently only the state of refresh blackhole is displayed.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the state of the ACL refresh blackhole, set by the `enable access-list refresh blackhole` command.

Example

The following command displays state of the ACL configuration:

```
show access-list configuration
```

The output from this command is similar to:

```
Access-list Refresh Blackhole: Enabled
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show access-list counter

```
show access-list counter {<countername>} {any | ports <portlist> | vlan
<vlanname>} {ingress | egress}
```

Description

Displays the specified access list counters.

Syntax Description

countername	Specifies the ACL counter to display.
portlist	Specifies to display the counters on these ports.
vlanname	Specifies to display the counters on the VLAN.
ingress	Specifies to display ingress counters.
egress	Specifies to display egress counters (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default direction is ingress.

Usage Guidelines

Use this command to display the ACL counters.

Example

The following example displays all the counters for all ACLs:

```
show access-list counter
```

The output of this command is similar to the following:

Policy Name	Vlan Name	Port	Direction	Byte Count
Counter Name		Packet Count		
=====				
firewall1	*	*	ingress	
DENY_SYN		0		0
PERMIT_SYN_ACK		1228300404		1920048848
tc_2.4	*	2:1	ingress	
arp192		3		204
denyAll		0		0
destIp		0		0
destIp2		0		0
tcp	*	2:7	ingress	
PERMIT_SYN_ACK		0		0


```

TCP_flag      *          3:6    ingress
  denyAll      0          0
  ipArp        0          0
  tcpflags-syn 0          0
qos_hongkong   *          3:8    ingress
  qp2cnt       0          0
  qp4cnt       0          0
  qp5cnt       0          0
tcp            v1         *      ingress
  PERMIT_SYN_ACK 3759119344 2217044928

```

The following example displays all the counters for the ACL on port 2:1:

```
show access-list counter port 2:1
```

The output of this command is similar to the following:

Policy Name	Vlan Name	Port	Direction	
Counter Name		Packet	Count	Byte Count
=====				
tc_2.4	*	2:1	ingress	
arp192		3		204
denyAll		0		0
destIp		0		0
destIp2		0		0
History				

History

This command was available in ExtremeXOS 10.1.

The egress options was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

show access-list dynamic

```
show access-list dynamic
```

Description

Displays the names of existing dynamic ACLs and a count of how many times each is used.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the names of existing dynamic ACLs, and how many times the ACL is used (bound to an interface).

To see the conditions and actions for a dynamic ACL, use the following command:

```
show access-list dynamic rule <rule> {detail}
```

Example

The following command displays names of all the dynamic ACLs:

```
show access-list dynamic
```

The following is sample output for this command:

```
Dynamic Rules:
Udpacl          Bound to 1 interfaces
icmp-echo       Bound to 1 interfaces
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show access-list dynamic counter

```
show access-list dynamic counter {{<countername>}} any
| {{<countername>}} ports <portlist>
| {{<countername>}} vlan <vlanname>}
{ingress | egress}
```

Description

Displays the dynamic ACL counters.

Syntax Description

countername	Display the counter.
any	Specifies the wildcard ACL.
portlist	Specifies which ports' ACLs to display.
vlanname	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default is to display all interfaces, ingress.

Usage Guidelines

None.

Example

The following command displays all the dynamic ACL counters:

```
show access-list dynamic counter
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show access-list dynamic rule

```
show access-list dynamic rule <rule> {detail}
```

Description

Displays the syntax of a dynamic ACL.

Syntax Description

rule	Specifies the rule to display.
detail	Specifies to display where the ACL has been applied.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the syntax of the dynamic ACL *udpACL*:

```
show access-list dynamic rule udpACL
```

The output of the command is similar to the following:

```
entry udpACL {
if match all {
    source-address 10.203.134.0/24 ;
    destination-address 140.158.18.16/32 ;
    protocol udp ;
    source-port 190 ;
    destination-port 1200 - 1250 ;
} then {
    permit ;
} }
```

The following command displays where the dynamic ACL *udpACL* has been applied:

```
show access-list dynamic rule udpACL
```

The output of the command is similar to the following:

```
Rule updacl has been applied to the following interfaces.
Vlan Name      Port      Direction
=====
*              1        ingress
```

History

This command was first available in ExtremeXOS 11.3.

The `detail` keyword was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show access-list interface

```
show access-list {rule <rule> {<start>} } [ any | port <port> | vlan
<vlanname> ] {zone <zone_name> { appl-name <appl_name> {priority <number>
}} } {ingress | egress} {detail}
```

Description

Displays the specified ACL zones, including their priority, applications, and the application priorities.

Syntax Description

any	Displays all zones on the specified interface.
vlan <vlanname>	Displays all ACLs associated with the specified VLAN.
port <port>	Displays all ACLs associated with the specified ports.
zone <zone_name>	Specifies a zone to be displayed.
appl-name <appl_name>	Displays information by application within a zone.
priority <number>	Displays ACLs of the specified priority only, within an application area.
ingress	Displays ACLs applied to traffic in the ingress direction.
egress	Displays ACLs applied to traffic in the egress direction.
detail	Displays all ACLs applied to the specified interface.

Default

N/A.

Usage Guidelines

Use this command to display the ACL zones, applications, and priorities.

Specifying a zone will show all the ACLs installed in the particular zone. Specifying a priority within a zone will show all the ACLs installed at a particular priority within a zone.

Use the `detail` keyword to display all ACLs installed on a given interface.

Example

The following example displays the detailed view of the ACLs on port 1:1

```
show access-list port 1:1 detail
```

The output of this command is similar to the following:

```
* BD-PC.1 # show access-list port 1:1 detail
RuleNo  Application    Zone          Sub Zone
=====
      1  CLI    myZone1
entry mac1 {
```

```

if match all {
    ethernet-source-address 00:0c:29:e5:94:c1 ;
    destination-address 192.168.11.144/32 ;
} then {
    count mac1 ;
} }

2 CLI myZone5
entry mac51 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
    count mack51;
} }

3 CLI myZone5
entry mac52 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
    count mac52 ;
} }

```

The following example displays the detailed view of the priority 5 ACLs in the zone *myzone* on port 1:1:

```

* BD-PC.2 # show access-list port 1:1 zone myZone priority 5 detail
RuleNo  Application      Zone      Sub Zone
=====
2 CLI myZone5
entry mac51 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
    count mack51;
} }

3 CLI myZone5
entry mac52 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
    count mac52 ;
} }

```

The following example displays the priority 5 ACLs in the zone *myzone* on port 1:1:

```

BD-PC.2 # show access-list port 1:1 zone myZone priority 5

#Dynamic Entries ((*)- Rule is non-permanent )
RuleNo      Name          Application      Zone      Sub-Zone
1           mac51         CLI             myZone     5
2           mac52         CLI             myZone     5

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all switches.

show access-list meter

```
show access-list counter {<metername>} [any | ports <portlist> | vlan
<vlanname>]
```

Description

Displays the specified access list meter statistics and configurations.

Syntax Description

metername	Specifies the ACL meter to display.
portlist	Specifies to display the meters on these ports.
vlanname	Specifies to display the meters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the ACL meters.

Example

The following example displays all the counters the ACL on port 7:1

```
show access-list meter port 2:1
```

The output of this command is similar to the following:

ACL Name	Vlan	Committed Rate (Kbps)	Port Max Burst Size (K)	Out-of-Profile Action	Out-of-Profile DSCP	Out-of-Profile Packet Count
donacl	*		7:1			
m1		10000	Max	Drop		0

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 family, SummitStack, and Summit family switches.

show access-list usage acl-mask

```
show access-list usage acl-mask port <port>
```

Description

Displays the number of ACL masks consumed by the ACLs on a particular port.

Syntax Description

port	Specifies to display the usage on this port.

Default

N/A.

Usage Guidelines

The BlackDiamond 8800 original series modules and Summit X450 series switches (whether or not included in a SummitStack) have a total of sixteen ACL masks per port on the switch. To avoid exhausting the masks available on the switch, you must carefully plan your use of ACL masks.

Use this command to display how many masks are currently consumed on a port.

Example

The following example displays the ACL mask usage on port 1:1:

```
show access-list usage acl-mask port 1:1
```

The output of this command is similar to the following:

```
Used: 3   Available: 12
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 8800 family, SummitStack, and Summit family switches.

show access-list usage acl-range

```
show access-list usage acl-range port <port>
```

Description

Displays the number of Layer 4 port ranges consumed by the ACLs on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage for the slices that support this port.

Default

N/A.

Usage Guidelines

The BlackDiamond 8800 a-series and e-series modules, SummitStack, and Summit X450a and X450e series switches can support a total of 16 Layer 4 port ranges among the slices that support each group of 24 ports.

Use this command to display how many of these Layer 4 ranges are currently consumed by the ACLs on the slices that support a particular port. The output of this command also displays which ports share the same slices as the specified port.

This command will return an error if the port specified is on a BlackDiamond 8800 original series module, or a Summit X450 series switch (whether or not included in a SummitStack).

Example

The following example displays the Layer 4 range usage on port 9:1:

```
show access-list usage acl-range port 9:1
```

The output of this command is similar to the following:

```
Ports 9:1-9:12, 9:25-9:36
L4 Port Ranges:  Used: 4  Available: 12
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 8800 a-series and e-series modules, SummitStack, and Summit X450a and X450e series switches.

show access-list usage acl-rule

```
show access-list usage acl-rule port <port>
```

Description

Displays the number of ACL rules consumed by the ACLs on a particular port or on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage on this port.

Default

N/A.

Usage Guidelines

The BlackDiamond 8800 original series modules and Summit X450 series switches (whether or not included in a SummitStack) support 128 rules per Gigabit Ethernet port and 1024 rules per 10 Gigabit Ethernet port, so you may need to plan your use of ACL rules. On these platforms, this command displays the rule usage of the port.

The BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches support 128 rules per slice for the slices that support each group of 24 ports. Since the BlackDiamond 8800 a-series modules and Summit X450a series switches support sixteen slices per 24-port group, there can be a maximum of 2048 rules supported for the group. Since the BlackDiamond 8800 e-series modules and Summit X450e series switches support 8 slices per 24-port group, there can be a maximum of 1024 rules supported for the group. On these platforms the command displays the rules used per slice, and also displays which ports share the same slices as the specified port.

Example

The following example displays the ACL rule usage on port 5:

```
show access-list usage acl-rule port 5
```

The output of this command on a Summit X450 series switch is similar to the following:

```
Used: 10   Available: 118
```

The output of this command on a Summit X450e series switch is similar to the following:

```
Ports 1-12, 25-36
Total Rules:      Used: 34   Available: 990
```

History

This command was first available in ExtremeXOS 11.4.

This command was modified to support BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e switches (whether or not included in a SummitStack) in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 8800 series switch, SummitStack, and Summit family switch.

show access-list usage acl-slice

```
show access-list usage acl-slice port <port>
```

Description

Displays the number of ACL slices and rules consumed by the ACLs on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage for the slices that support this port.

Default

N/A.

Usage Guidelines

The BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches (whether or not included in a SummitStack) support 128 rules per slice for the slices that support each group of 24 ports. Since the BlackDiamond 8800 a-series modules and Summit X450a series switches support sixteen slices per 24-port group, there can be a maximum of 2048 rules supported for the group. Since the BlackDiamond 8800 e-series modules and Summit X450e series switches support 8 slices per 24-port group, there can be a maximum of 1024 rules supported for the group.

Use this command to display how many slices, and how many rules per each slice are currently consumed by the ACLs on the slices that support a particular port. This command also displays which ports share the same slices as the specified port.

Example

The following example displays the ACL slice usage on port 8:1:

```
show access-list usage acl-slice port 8:1
```

The output of this command is similar to the following:

```
Ports 8:1-8:12, 8:25-8:36
Slices:      Used: 3  Available: 5
Slice 5 Rules:  Used: 9  Available: 119
Slice 6 Rules:  Used: 1  Available: 127
Slice 7 Rules:  Used: 24  Available: 104
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches (whether or not included in a SummitStack).

unconfigure access-list

```
unconfigure access-list {any | ports <portlist> | vlan <vlanname>} {ingress
| egress}
```

Description

Removes a policy file ACL from the specified interface.

Syntax Description

aclname	Specifies the ACL name. The name can be from 1-32 characters long.
portlist	Specifies the ports on which this ACL is applied.
vlanname	Specifies the VLAN on which this ACL is applied.
ingress	Remove the ACL for packets entering the switch on this interface.
egress	Remove the ACL for packets leaving the switch from this interface (BlackDiamond 10808 and BlackDiamond 12800 series only).

Default

The default direction is ingress.

Usage Guidelines

This command removes ACLs that are contained in ACL policy files. To remove dynamic ACLs, use the following command:

```
configure access-list delete <dynamic_rule> [any | vlan <vlanname> | ports <portlist>
| all] {ingress | egress} {application <appl_name>}
```

To remove all non-dynamic ACLs from all interfaces, do not specify any ports or VLANs.

Example

The following command removes the ACL from port 1:2:

```
unconfigure access-list ports 1:2
```

The following command removes the ACLs from ports 1:2-6:3 and 7:1:

```
unconfigure access-list ports 1:2-6:3,7:1
```

The following command removes the wildcard ACL:

```
unconfigure access-list any
```

The following command removes all ACLs from all the interfaces, including the wildcard ACL:

```
unconfigure access-list
```

History

This command was available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress options was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms. The egress option is available on BlackDiamond 10808 and BlackDiamond 12800 series switches only.

18 QoS Commands

This chapter describes commands for:

- Configuring Quality of Service (QoS) profiles
- Creating traffic groupings and assigning the groups to QoS profiles
- Configuring, enabling and disabling explicit class-of-service traffic groupings (802.1p and DiffServ)
- Configuring traffic grouping priorities
- Metering using ACLs—BlackDiamond 8800 series of switches, SummitStack, and Summit family of switches only
- Verifying configuration and performance
- Configuring and displaying bidirectional rate-shaping—BlackDiamond 10808s switch only
- Egress traffic rate limiting—BlackDiamond 8800 series switches, SummitStack, and Summit family of switches only
- Hierarchical QoS—BlackDiamond 12800 R-series switch only
 - Strict priority mode
 - Bandwidth mode—meets the specifications of the Metro Ethernet Forum (MEF) 14 standard

Quality of Service (QoS) is a feature of ExtremeXOS that allows you to specify different service levels for outbound and inbound traffic. QoS is an effective control mechanism for networks that have heterogeneous traffic patterns. Using QoS, you can specify the service that a traffic type receives.

Policy-based QoS allows you to protect bandwidth for important categories of applications or specifically limit the bandwidth associated with less critical traffic. The switch contains separate hardware queues on every physical port. Each hardware queue is programmed by ExtremeXOS with bandwidth management and prioritization parameters, defined as a QoS profile. The bandwidth management and prioritization parameters that modify the forwarding behavior of the switch affect how the switch transmits traffic for a given hardware queue on a physical port. Up to eight physical queues per port are available.

With the 11.0 version of the software, you can also assign parameters to traffic ingressing the switch for minimum and maximum bandwidth.

To configure QoS, you define how your switch responds to different categories of traffic by creating and configuring QoS profiles. The service that a particular type of traffic receives is determined by assigning a QoS profile to a traffic grouping or classification. The building blocks are defined as follows:

- QoS profile—Defines bandwidth and prioritization parameters.
- Traffic grouping—A method of classifying or grouping traffic that has one or more attributes in common.
- QoS policy—The combination that results from assigning a QoS profile to a traffic grouping.

QoS profiles are assigned to traffic groupings to modify switch-forwarding behavior. When assigned to a traffic grouping, the combination of the traffic grouping and the QoS profile comprise an example of a single policy that is part of policy-based QoS.

Extreme Networks switch products support explicit Class of Service traffic groupings. This category of traffic groupings describes what is sometimes referred to as *explicit packet marking*, and includes:

- IP DiffServ code points, formerly known as IP TOS bits
- Prioritization bits used in IEEE 802.1p packets

All Extreme Networks switches support the standard 802.1p priority bits that are part of a tagged Ethernet packet.

With the 11.0 version of the software on the BlackDiamond 10808 switch, you can also assign parameters to traffic ingressing the switch for minimum and maximum bandwidth and priority queuing to the backplane.

Beginning with ExtremeXOS version 11.4, Extreme Networks introduces hardware-based hierarchical QoS (HQoS), or rate limiting, on the BlackDiamond 12800 R-series switch in strict priority mode. ExtremeXOS version 12.0 adds the bandwidth mode, or MEF-14 compatible mode, which allows you to set committed and peak rates.

clear counters traffic queue

```
clear counters traffic queue [all | <queue_name>]
```

Description

Clears all counters used in gathering statistics for HQoS.

Syntax Description

all	Clears HQoS counters for all traffic queues on the switch.
queue_name	Enter the name of the traffic queue on which you want to clear all HQoS counters.

Default

N/A.

Usage Guidelines

Use this command to clear the HQoS counters on a specified traffic queue or on all traffic queues.



NOTE

If you issue the command `clear counters` with no keywords, the system clears all the application counters on the software.

Example

The following command clears the HQoS counters for all traffic queues on the switch:

```
clear counters traffic queue all
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

configure diffserv examination code-point

```
configure diffserv examination code-point <code-point> {qosprofile}
<qosprofile>
```

Description

Configures the default ingress DiffServ code points (DSCP) to QoS profile mapping.

Syntax Description

code-point	Specifies a DiffServ code point (a 6-bit value in the IP-TOS byte in the IP header).
qosprofile	Specifies the QoS profile to which the DiffServ code point is mapped.

Default

See [Table 21](#).

Usage Guidelines

You can specify up to 64 different code points for each port. Code point values are grouped and assigned to the default QoS profiles as shown in [Table 21](#).

Table 21: Default DiffServ code point-to-QoS profile mapping

Code point	BlackDiamond 10808 and 12800 series switches QoS profile	BlackDiamond 8800 series switches, SummitStack, and Summit family of switches QoS profile
0-7	QP1	QP1
8-15	QP2	QP1
16-23	QP3	QP1
24-31	QP4	QP1
32-39	QP5	QP1
40-47	QP6	QP1
48-55	QP7	QP1
56-63	QP8	QP8

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot set per port DiffServ examination code points. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure examining DiffServ information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.

Example

The following command specifies that code point 25 be assigned to QP2:

```
configure diffserv examination code-point 25 qosprofile qp2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure diffserv replacement

```
configure diffserv replacement [{qosprofile} <qosprofile> | priority
<value>] code-point <code_point>
```

Description

Configures the default egress DiffServ replacement mapping.

Syntax Description

qosprofile	Specifies a QoS profile.
value	Specifies the 802.1p priority value.
code_point	Specifies a 6-bit value to be used as the replacement code point in the DiffServ (IP-TOS byte) of the IP header.

Default

N/A.

Usage Guidelines



NOTE

Extreme Networks recommends that you use the qosprofile <qosprofile> value to configure this parameter.

The default QoS profile to 802.1p priority value to code point mapping is shown in [Table 22](#).

Table 22: Default QoS profile-to-802.1p priority value-to-code point

802.1p priority value	BlackDiamond 10808 and 12800 series switches QoS profile	BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches QoS profile	Code point
0	QP1	QP1	0
1	QP2	QP1	8
2	QP3	QP1	16
3	QP4	QP1	24
4	QP5	QP1	32
5	QP6	QP1	40
6	QP7	QP1	48
7	QP8	QP8	56

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ replacement feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure replacing DiffServ information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.

Example

The following command specifies that a code point value of 5 should be used to replace the DiffServ (TOS) bits in packets in QP2:

```
configure diffserv replacement qosprofile qp2 code-point 5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure dot1p type

```
configure dot1p type <dot1p_priority> {qosprofile} <qosprofile>
```

Description

Configures the default QoS profile to 802.1p priority mapping.

Syntax Description

dot1p_priority	Specifies the 802.1p priority value. The value is an integer between 0 and 7.
qosprofile	Specifies a specific QoS profile.

Default

The default mapping of each 802.1p priority value to QoS profile is shown in [Table 23](#).

Table 23: Default 802.1p priority value-to-QoS profile mapping

Priority value	BlackDiamond 10808 and 12800 series switches QoS profile	BlackDiamond 8800 series switches, SummitStack, and Summit family of switches default QoS profile
0	QP1	QP1
1	QP2	QP1
2	QP3	QP1
3	QP4	QP1
4	QP5	QP1
5	QP6	QP1
6	QP7	QP1
7	QP8	QP8

Usage Guidelines

An 802.1p priority value seen on ingress can be mapped to a particular QoS profile and with specific bandwidth management and priority behavior.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the 802.1p feature. If you attempt to configure 802.1p information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.

BlackDiamond 8800 series switches, SummitStack, and Summit family of switches only. You must create the QoS profile first, using the `create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]` command, to map the 802.1p information to QoS profile 2 through 7.

SummitStack only. You must create the QoS profile first, using create qosprofile [QP2 | QP3 | QP4 | QP5] command, to map the 802.1p information to QoS profile 2 through 5. You cannot create QP6 or QP7 in a SummitStack.

Example

The following commands reassign (from the default) the QoS profiles associated with 802.1p priority values 1 and 2:

```
configure dot1p type 2 qosprofile qp2  
configure dot1p type 1 qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure meter

```
configure meter <metername> {max-burst-size <burst-size> [Gb | Kb | Mb]}
{committed-rate <cir> [Gbps | Mbps | Kbps]} {out-actions [drop | set-drop-
precedence {dscp [none | <dscp-value>]}]}
```

Description

Configures an ACL meter.

Syntax Description

metername	Specifies the ACL meter name.
max-burst-size	Specifies the maximum burst size.
committed-rate	Specifies the committed rate.
out-actions	Specifies actions to take if traffic exceeds the profile.
drop	Specifies to drop out of profile traffic.
set-drop-precedence	Specifies to mark packet for high drop precedence.
dscp	Specifies to set DSCP.
none	Specifies to leave the DSCP value unchanged.

Default

By default, a newly committed meter has no maximum burst size, no committed rate, and a default action of drop.

Usage Guidelines

The ACL meter configured with this command will be associated with an ACL by specifying the meter name in the action modifier field of the ACL entry.

The `committed-rate` keyword specifies the traffic rate allowed for this meter. The committed traffic rate may be specified in Kbps, Mbps, or Gbps. The `committed-rate` parameter range is 64Kbps to 1 Gbps for GE ports and 1 Mbps to 10 Gbps for 10GE ports. The `max-burst-size` keyword specifies the maximum number of consecutive bits that are allowed to be in-profile at wire-speed. The `burst-size` parameter may be specified in Kb, Mb, or Gb. The specified `burst-size` will be rounded down to the nearest supported size. For GE ports, the burst sizes are rounded down to the nearest 32Kb, 64Kb, 128Kb, 256Kb, 512Kb, 1Mb, 2Mb, 4Mb, and additionally 8Mb, 16Mb, 32Mb, 64Mb, 128Mb, 256Mb, 512Mb, or 1Gb for 10 GE ports.

The keyword `out-actions` specifies the action that is taken when a packet is out-of-profile. The supported actions include dropping the packet, marking the drop precedence for the packet, or setting the DSCP value in the packet. The keyword `drop` indicates that any out-of-profile packet is immediately dropped. The keyword `set-drop-precedence` marks out-of-profile packets with high drop precedence. If the optional keyword `set-dscp` is specified, the DSCP value, as specified by the parameter `dscp-value`, is written into the out-of-profile packet. Setting the DSCP value to `none` leaves the DSCP value in the packet unchanged.

Example

The following command configures the ACL meter *maximum_bandwidth*, assigns it a rate of 10 Mbps, and sets the out of profile action to drop:

```
configure meter maximum_bandwidth committed-rate 10 Mbps out-action drop
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure meter committed-rate peak-rate

```
configure meter <metername> [committed-rate <cir> [Gbps | Mbps | Kbps] |  
peak-rate <pr>[Gbps | Mbps | Kbps]]
```

Description

Configures a peak rate for a meter for use in bandwidth mode HQoS.

Syntax Description

metername	Specifies the meter name.
cir	Specifies the committed rate size.
pr	Specifies the peak rate size.

Default

By default, a newly committed meter has no committed or peak rate.

Usage Guidelines



NOTE

This command is available only in bandwidth mode HQoS; it is not available in strict priority mode HQoS.

You associate the meter configured with this command with an ACL by specifying the meter name.

For bandwidth HQoS, enter the committed rate and the peak rate allowed for each meter. You can configure and associate a different meter on each of the four sub-queues for each ingress traffic queue that allows egress shaping associated with this ACL. Each of the sub-queues corresponds to the L1 queues or the four CoS value pairs, as follows:

- COS0-1
- COS2-3
- COS4-5
- COS6-7

You can apply meters to each CoS value pair independently or to the entire traffic queue (by specifying all CoS pairs).



NOTE

You must configure the peak rate higher than the committed rate.

Example

The following command configures the meter *bandwidth_rate* and assigns it a committed rate of 5 Mbps and a peak rate of 10 Mbps:

```
configure meter bandwidth_rate committed-rate 5 Mbps peak-rate 10 Mbps
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch with an enabled MEF Feature Pack.

configure meter peak-rate

```
configure meter <metername> [peak-rate <pr> [Gbps | Mbps | Kbps]]
```

Description

Configures a peak rate for a meter for use in strict priority mode HQoS.

Syntax Description

metername	Specifies the meter name.
pr	Specifies the peak rate size.

Default

By default, a newly committed meter has no peak rate.

Usage Guidelines



NOTE

Use this command to configure the meters for strict priority mode HQoS; for configuring meters in bandwidth mode HQoS, use the [configure meter <metername> \[committed-rate <cir> \[Gbps | Mbps | Kbps\] | peak-rate <pr>\[Gbps | Mbps | Kbps\]\]](#) command.

You associate the meter configured with this command with an ACL by specifying the meter name.

For strict priority HQoS, enter the peak rate allowed on the ingress queue that allows egress shaping associated with this ACL.

Example

The following command configures the meter *maximum_bandwidth* and assigns it a peak rate of 10 Mbps:

```
configure meter maximum_bandwidth peak-rate 10 Mbps
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

configure ports qosprofile

```
configure ports <port_list> {qosprofile} <qosprofile>
```

Description

Configures one or more ports to use a particular egress QoS profile.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
qosprofile	Specifies a QoS profile.

Default

All ports have the default qosprofile of QP1.

Usage Guidelines

Extreme switches support eight QoS profiles (QP1 to QP8) for each port. SummitStack does not permit configuration of either QP6 or QP7.



NOTE

This command applies only to untagged packets on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. You must use this command to assign a QoS profile to ports. This command applies to egress QoS profiles only; the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches support only egress QoS profiles.

BlackDiamond 10808 and 12800 series switches. This command refers to egress QoS profiles; it assigns traffic ingressing this port to a specified egress QoS profile. Use the `configure ports rate-limit map cos high` command to assign the ingress QoS profiles to ports; the ingress QoS profiles assign ingressing traffic to a specified queue to the backplane. This command applies only to untagged packets on the BlackDiamond 10808 and 12800 series switches.

Example

The following command configures port 5 on slot 5 on a modular switch to use QoS profile QP3:

```
configure ports 5:5 qosprofile QP3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure ports rate-limit egress

```
configure ports <port_list> rate-limit egress [no-limit | <cir-rate> [Kbps
| Mbps | Gbps] {max-burst-size <burst-size> [Kb | Mb]]]
```

Description

Configures an egress traffic rate limit for a port or groups of ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-limit	Specifies traffic be transmitted without limit; use to reconfigure or unconfigure previous rate-limiting parameters.
cir-rate	Specifies the desired rate limit in Kbps, Mbps, or Gbps.
max-burst-size	Specifies amount of traffic above the cir-rate that is allowed to burst (for a short duration) from the port in K bits (Kb) or M bits (Mb).

Default

No-limit.

Usage Guidelines

Port speed limits the egress traffic, as follows:

- 1 Gbps port—64 Kbps increments
- 10 Gbps port—1 Mbps increments

If the specified egress limit (`cir-rate`) is not a multiple of 64 Kbps for a 1 Gbps port or 1 Mbps for a 10 Gbps port, the specified value is rounded down to the nearest appropriate multiple based on the port type.

Use the `no-limit` parameter to:

- Unconfigure egress rate limiting on the port(s)
- Reconfigure existing egress rate limiting on the port(s)

The `max-burst-size` parameter is the amount of traffic above the value in the `cir-rate` parameter that is allowed to burst from the port(s) for a short duration.

Example

The following command configures egress rate-limiting on slot 3 port 1 on a modular switch for 3 Mbps and a maximum burst size of 5 M bits:

```
configure port 3:1 rate-limit egress 3 Mbps max-burst-size 5 Mb
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure ports rate-limit egress aggregate

```
configure ports <port_list> rate-limit egress aggregate [{committed_rate
<committed_bps> [k | m ]} {peak_rate <peak_bps> [k | m]}} | {{minbw
<minbw_number>} {maxbw <maxbw_number>}}]
```

Description

Sets the aggregate egress committed rate and peak rate or minimum bandwidth and maximum bandwidth on the specified ports.

Syntax Description

ports	Specifies the port or ports to configure port-based egress rate limiting on.
committed_rate	Specifies an absolute number in K bits or M bits for a guaranteed committed rate and for a minimum rate for this egress port.
peak_rate	Specifies an absolute number in K bits or M bits for a maximum rate for this egress port.
minbw	Specifies the minimum egress percentage bandwidth allowed for specified ports.
maxbw	Specifies the maximum egress percentage bandwidth allowed for specified ports.

Default

Minimum is 0%; maximum is 100%.

Usage Guidelines

Sets the aggregate committed rate and peak rate or the minimum bandwidth, and maximum bandwidth on the specified ports.

Ensure that the value entered for the minimum bandwidth is *less* than the value entered for the maximum bandwidth and the value for committed rate is *less* than the value entered for peak rate.

To remove the rate limit on egress bandwidth on a port, re-issue this command using the default values.

To display the configured committed rate, peak rate, minimum bandwidth, and maximum bandwidth for the specified ports, issue the `show ports <port_list> information detail` command.

The following is sample output for this command (the Aggregate Queue information displays the settings for this command):

```
Port:      2:5
Virtual-router: VR-Default
Type:      NONE
Random Early drop:      Unsupported
Admin state:      Enabled with auto-speed sensing auto-duplex
Link State:      Active, 1Gbps, full-duplex
Link Counter: Up      0 time(s)
VLAN cfg:
```

```

STP cfg:

Protocol:
Trunking:      Load sharing is not enabled.
EDP:           Enabled
ELSM:          Disabled
DLCS:          Unsupported
lbDetect:      Unsupported
Learning:      Enabled
Unicast Flooding:  Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo:         Disabled
BG QoS monitor: Unsupported
High Priority COS: 7
Packet Length Adjustment: 0 bytes
QoS Profile:    None configured
Aggregate Queue:
    QP0  MinBw =          0% MaxBw =          75%  Pri = 8
Queue:
    QP1  MinBw =          0% MaxBw =         100%  Pri = 1
    QP2  MinBw =          0% MaxBw =         100%  Pri = 2
    QP3  MinBw =          0% MaxBw =         100%  Pri = 3
    QP4  MinBw =          0% MaxBw =         100%  Pri = 4
    QP5  MinBw =          0% MaxBw =         100%  Pri = 5
    QP6  MinBw =          0% MaxBw =         100%  Pri = 6
    QP7  MinBw =          0% MaxBw =         100%  Pri = 7
    QP8  MinBw =          0% MaxBw =         100%  Pri = 8
Ingress Rate Shaping :      Unsupported
Ingress IPTOS Examination:  Disabled
Egress IPTOS Replacement:   Disabled
Egress 802.1p Replacement:  Disabled
NetLogin:                   Disabled
NetLogin port mode:         Port based VLANs
Smart redundancy:           Enabled
Software redundant port:    Disabled
Preferred medium:           Fiber

```

Example

The following command configures the egress bandwidth on port 2:5 of the BlackDiamond 12800 R-series switch as a maximum of 75% of the total bandwidth:

```
configure ports 2:5 rate-limit egress aggregate maxbw 75
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

The command is available only on the BlackDiamond 10808 and 12800 series switches.

configure ports rate-limit map cos high

```
configure ports <port_list> rate-limit map cos high <cos_value>
```

Description

Maps the CoS value to the high queue at ingress in strict priority HQoS. In this HQoS mode, the system has three queues: high, medium, and low. Using the ACLs to classify traffic, the system assigns packets to one of these queues based on the CoS value, with the other CoS values automatically cascading the packets to the other queues.

Syntax Description

ports	Specifies the port or ports you are mapping the high queue CoS value for.
cos_value	Specifies the Cos value you want for the high queue for the specified ports. You can configure this value from 7 to 2.

Default

The default CoS value to queue mapping on ingressing packets is as follows:

- High—7
- Medium—6
- Low—5, 4, 3, 2, 1, 0, and untagged packets

Usage Guidelines



NOTE

This command is available only in strict priority HQoS mode; it is not available in bandwidth mode HQoS.

With strict priority mod HQoS, you can configure which CoS value directs the packet to the high queue, and the remaining CoS values are automatically queued. You can configure this value from 7 to 2.

If you attempt to enter a value other than 2 to 7, the system returns the following error message:

```
Error: highPriCosVal value can range between 2 and 7
```



NOTE

All packets with a CoS value higher than that configured for the high queue automatically go to the low queue. (If you configure 2 as the CoS value for the high queue, packets with a CoS value of 3 to 7 all go to the low queue.)

For example, if you configure 5 as the CoS value for the high queue, the queues contain packets with the following CoS values:

- High contains 5
- Medium contains 4
- Low contains 3, 2, 1, 0, 7, 6, and those with no CoS value

To display the configured CoS value mapped to the high queue for the specified ports for strict priority rate limiting, issue the `show port information detail` command (the setting displays on the High Priority COS line).

**NOTE**

You must use the detail variable to display this information.

The following is sample output from the `show port information detail` command:

```
show port 2:1 information detail
Port: 2:1
    Virtual-router: VR-Default
    Type: UTP
    Random Early drop: Unsupported
    Admin state: Enabled with auto-speed sensing auto-duplex
    Link State: Active, 1Gbps, full-duplex
    Link Counter: Up 8 time(s)
    VLAN cfg:
        Name: vman110, Internal Tag = 110, MAC-limit = No-limit, Virtual
router: VR-Default

    STP cfg:

    Protocol:
        Name: vman110 Protocol: ANY Match all protocols.
    Trunking: Load sharing is not enabled.
    EDP: Enabled
    ELSM: Disabled
    Learning: Enabled
    Unicast Flooding: Enabled
    Multicast Flooding: Enabled
    Broadcast Flooding: Enabled
    Jumbo: Enabled, MTU= 9216
    High Priority COS: 7
    Packet Length Adjustment: 0 bytes
    QoS Profile: None configured
    Aggregate Queue:
        QP0 MinBw = 0% MaxBw = 75% Pri = 8
    Queue:
        QP1 MinBw = 0% MaxBw = 100% Pri = 1
        QP2 MinBw = 0% MaxBw = 100% Pri = 2
        QP3 MinBw = 0% MaxBw = 100% Pri = 3
        QP4 MinBw = 0% MaxBw = 100% Pri = 4
        QP5 MinBw = 0% MaxBw = 100% Pri = 5
        QP6 MinBw = 0% MaxBw = 100% Pri = 6
        QP7 MinBw = 0% MaxBw = 100% Pri = 7
        QP8 MinBw = 0% MaxBw = 100% Pri = 8
```


Ingress Rate Shaping :	Unsupported
Ingress IPTOS Examination:	Disabled
Egress IPTOS Replacement:	Disabled
Egress 802.1p Replacement:	Disabled
NetLogin:	Disabled
NetLogin port mode:	Port based VLANs
Smart redundancy:	Enabled
Software redundant port:	Disabled
Preferred medium:	Fiber

Example

The following command sets the high queue to a CoS value of 6 on the ingressing packet on ports 4:2 to 4:10 (the other values follow in cascading order):

```
configure ports 4:2-10 rate-limit map cos high 6
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

The command is available only in strict priority HQoS mode on the BlackDiamond 12800 R-series switch.

configure ports rate-limit packet byte-adjustment

```
configure ports <port_list> rate-limit packet byte-adjustment [increase
<add_bytes> | decrease <sub_bytes>]
```

Description

Configures the number of bytes to add to or subtract from each packet on the specified ports for calculating the ingress traffic rate, traffic utilization, and traffic statistics for HQoS.

Syntax Description

port_list	Specifies the port or ports that will have the specified bytes added to or subtracted from all packets in calculating the ingress traffic rate, traffic utilization, and traffic statistics.
add_bytes	Specifies the number of bytes to add to each packet when calculating the ingress traffic rate. The default number is 0; the available range of added bytes is 1 to 4.
sub_bytes	Specifies the number of bytes to subtract from each packet when calculating the ingress traffic rate. The default number is 0; the available range of subtracted bytes is 1 to 4.

Default

0.

Usage Guidelines

With HQoS, use this command to specify the number of bytes to be added to or subtracted from all packets ingress the specified ports when calculating the ingress traffic rate, traffic utilization, and traffic statistics. You can add or subtract 1 to 4 bytes from each packet. To unconfigure this setting, re-issue the command and enter the value 0.

To display the number of bytes added to or subtracted from each ingress packet for calculating the ingress traffic rate, traffic utilization, and traffic statistics, use the `show ports <port_list> information detail` command.

The following is sample output from this command (the setting is displayed on the Packet Length Adjustment line):

```
show port 2:1 information detail
Port:      2:1
  Virtual-router: VR-Default
  Type:          UTP
  Random Early drop:      Unsupported
  Admin state:      Enabled with auto-speed sensing  auto-duplex
  Link State:      Active, 1Gbps, full-duplex
  Link Counter: Up      8 time(s)
  VLAN cfg:
    Name: vman110, Internal Tag = 110, MAC-limit = No-limit, Virtual
router: VR-Default
```

```

STP cfg:

Protocol:
    Name: vman110          Protocol: ANY          Match all protocols.
Trunking:      Load sharing is not enabled.
EDP:           Enabled
ELSM:          Disabled
Learning:      Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo: Enabled, MTU= 9216
High Priority COS: 7
Packet Length Adjustment: 0 bytes
QoS Profile:    None configured
Aggregate Queue:
    QP0  MinBw =          0% MaxBw =          75%  Pri = 8
Queue:
    QP1  MinBw =          0% MaxBw =          100%  Pri = 1
    QP2  MinBw =          0% MaxBw =          100%  Pri = 2
    QP3  MinBw =          0% MaxBw =          100%  Pri = 3
    QP4  MinBw =          0% MaxBw =          100%  Pri = 4
    QP5  MinBw =          0% MaxBw =          100%  Pri = 5
    QP6  MinBw =          0% MaxBw =          100%  Pri = 6
    QP7  MinBw =          0% MaxBw =          100%  Pri = 7
    QP8  MinBw =          0% MaxBw =          100%  Pri = 8
Ingress Rate Shaping :      Unsupported
Ingress IPTOS Examination:  Disabled
Egress IPTOS Replacement:   Disabled
Egress 802.1p Replacement:  Disabled
NetLogin:                   Disabled
NetLogin port mode:         Port based VLANs
Smart redundancy:            Enabled
Software redundant port:     Disabled
Preferred medium:            Fiber

```

Example

The following command instructs the system to subtract 4 bytes from all packets ingressing ports 2:1 to 2:10 when calculating ingressing traffic rates, traffic utilization, and traffic statistics:

```
configure ports 2:1-10 rate-limit packet byte-adjustment decrease 4
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

The basic command is available only on the BlackDiamond 12800 R-series switch.

configure qosprofile ingress ports

```
configure qosprofile ingress <iqp> [{committed_rate <committed_bps>
[k | m]} {maxbw <maxbw_number>} {minbw <minbw_number>} {peak_rate
<peak_bps> [k | m] {priority [<priority> | <priority_number>}}] ports
[<port_list> | all]
```

Description

Sets the ingress rate shaping parameters, which is an ingress QoS profile.

Syntax Description

iqp	Specifies an ingress QoS profile: - for 1G I/O modules—iqp1 and iqp2 - for 10G I/O modules—iqp1 to iqp8
committed_rate	Specifies an absolute number in K bits or M bits for a guaranteed committed rate and for a minimum rate for this ingress queue to the backplane.
maxbw	Specifies maximum allowed input bandwidth as a percentage of port speed. The range is 0 to 100%, and the default value is 100. Percentages are based on actual or configured port speed.
minbw	Specifies a guaranteed minimum bandwidth for this ingress QoS queue as a percentage of port speed. The range is 0 to 100%, and the default value is 0. Cumulative percentages of the queue on a given port should not exceed 100%. Percentages are based on actual or configured port speed.
peak_rate	Specifies an absolute number in K bits or M bits for a maximum rate for this ingress queue to the backplane.
priority	Specifies a service priority setting for the specified ingress queue to the backplane; this priority sets the priorities among the available ingress queues for the specified port.
priority_number	This priority sets the priorities among the available ingress queues for the specified port. The available queues vary according to the I/O module, as follows: 1G I/O module—2 queues and 2 priorities available; values are 1 (takes 1-4) or 2 (takes 5-8); or Low and LowHi. 10G module—8 queues and 8 priorities available; values are 1 to 8, with 8 being the highest priority, or Low, LowHi, Normal, NormalHi, Medium, MediumHi, High, and HighHi.
port_list	Specifies a list of slots and ports to which the parameters apply. May specify in the form 3-5, 2:5, 2:6-2:8.
all	Specifies this applies to all ports on the device.

Default

Disabled by default.

- Minimum bandwidth—0%
- Maximum bandwidth—100%

- **Priority**—By default, each qosprofile is assigned a different priority level, which varies by I/O module:
 - 1G I/O module:
 - IQP1 - 1, Low
 - IQP2 - 2, LowHi
 - 10G I/O module:
 - IQP1 - 1, Low
 - IQP2 - 2, LowHi
 - IQP3 - 3, Normal
 - IQP4 - 4, NormalHi
 - IQP5 - 5, Medium
 - IQP6 - 6, MediumHi
 - IQP7 - 7, High
 - IQP8 - 8, HighHi (highest priority)

Usage Guidelines

The number of ingress queues per port varies between the 1G I/O module and the 10G module.

On the 1G module, you have two ingress queues per port. The priority values of 1 to 4 map to the first queue, and the priority values of 5 to 8 map to the second queue.

On the 10G module, you have eight ingress queues per port. The priority values of 1 to 8 map one to each of the eight queues.

Example

The following command configures the ingress rate shaping parameters of QoS profile IQP3 for specified ports, using bandwidth percentages:

```
configure qosprofile ingress iqp3 minbw 27 maxbw 57 priority 4 ports 3:2
```

The following command configures the ingress rate shaping parameters for QoS profile IQP3 for all ports, using absolute values for committed rate and peak rate:

```
configure qosprofile ingress iqp3 committed-rate 64 k peak-rate 1000 k priority 4
ports all
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

The basic command is available only on the BlackDiamond 10808 switch.

configure qosprofile ports

```
configure qosprofile {egress} <qosprofile> [{committed_rate <committed_bps>
[k | m]} {maxbw <maxbw_number>} {minbw <minbw_number>} {peak_rate
<peak_bps> [k | m] {priority [<priority> | <priority_number>}}]
ports [<port_list> | all]
```

Description

Modifies the default egress QoS profile parameters.

Syntax Description

qosprofile	Specifies a QoS profile name. Range is QP1 to QP8.
committed_rate	Specifies an absolute number in K bits or M bits for a guaranteed committed rate and for a minimum rate for this egress queue.
maxbw	Specifies maximum allowed input bandwidth as a percentage of port speed. The range is 0 to 100%, and the default value is 100. Percentages are based on actual or configured port speed.
minbw	Specifies a guaranteed minimum bandwidth for this egress QoS queue as a percentage of port speed. The range is 0 to 100%, and the default value is 0. Cumulative percentages of the queue on a given port should not exceed 100%. Percentages are based on actual or configured port speed.
peak_rate	Specifies an absolute number in K bits or M bits for a maximum rate for this egress queue.
priority	Specifies a service priority setting for the specified egress queue; this priority sets the priorities among the available egress queues for the specified port.
priority_number	NOTE: This parameter applies only to the BlackDiamond 10808 and 12800 series switches. Specifies a service priority setting among the available egress ports for the specified port. Settings are: • 1, or Low • 2, or LowHi • 3, or Normal • 4 or NormalHi • 5, or Medium • 6, or MediumHi • 8, or HighHi This is the setting for the egress mode. NOTE: This parameter applies to the BlackDiamond 10808 and 12800 series switches only.
port_list	Specifies a list of slots and ports to which the parameters apply. May be in the form 3-5, 2:5, 2:6-2:8.
all	Specifies this applies to all ports on the device.

Default

- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Priority—By default, each qosprofile is assigned a different priority level:
 - QP1 - 1, Low (the lowest priority)
 - QP2 - 2, LowHi
 - QP3 - 3, Normal
 - QP4 - 4, NormalHi
 - QP5 - 5, Medium
 - QP6 - 6, MediumHi
 - QP7 - 7, High
 - QP8 - 8, HighHi (highest priority)

Usage Guidelines



NOTE

This command is not supported on BlackDiamond 8800 original modules.

The maximum bandwidth value can be configured as either:

- an absolute percentage of the total maximum link speed, regardless of the currently configured or negotiated speed
- OR
- an absolute peak rate in Mbps or Kbps

If you attempt to apply this command to ports on the original BlackDiamond 8800 modules, the system returns the following error:

```
ERROR: Setting minbw/maxbw is not supported on MSM-G8X I/O ports and G48T, G48P, G24X,
and 10G4X modules.
```

Example

The following command configures the egress QoS profile parameters of QoS profile Qp5 for specific ports on a BlackDiamond 10808 switch:

```
configure qosprofile qp5 minbw 10 maxbw 80 priority HighHi ports 5:5-5:7
```

History

This command was first available in ExtremeXOS 10.1.

Committed and peak rates were added in ExtremeXOS 11.0. Also in ExtremeXOS 11.0, ports were made mandatory.

Support for the BlackDiamond 8800 a-series and e-series modules and the Summit X450a and X450e series switches was added in ExtremeXOS 11.5.

Platform Availability

This command is available on the BlackDiamond 10808, Black Diamond 12800 series switches, BlackDiamond 8800 a-series and e-series modules, and the Summit X450a and X450e series switches, whether or not included in a SummitStack.

configure qosprofile weight

```
configure qosprofile <qosprofile> {maxbuffer <percent>} {weight <value>}
```

Description

Modifies the parameters for the default QoS profiles (QP1 and QP8) and defines the parameters for any user-created QoS profiles (QP2, QP3, QP4, QP5, QP6, and/or QP7). On a SummitStack, QP6 and QP7 are not available to the user.

Syntax Description

qosprofile	Specifies a QoS profile name. Range is QP1 to QP8; the defaults are QP1 and QP8.
maxbuffer <percent>	Specifies the percentage of the total buffer you are reserving for this QoS profile. The range is 1 to 100; the default setting is 100.
weight <value>	Specifies the weight value used for queue service weighting in the weighted-round-robin scheduler for this QoS profile. The range is 1 to 16; default is 1.

Default

- QoS profiles—QP1 and QP8
- Maximum buffer—100%
- Weight—1

Usage Guidelines

You must create a QoS profile before you can configure the user-created QoS profiles. You can modify the settings on the default QoS profiles (QP1 and QP8) because they are already created.

To assign a QoS profile to a port, use the `configure ports <port_list> {qosprofile} <qosprofile>` command. If you attempt to configure a QoS profile that you have not created, the system returns an error message. To configure the type of scheduling you want to use for the entire switch, use the `configure qosscheduler [strict-priority | weighted-round-robin]` command.

The `maxbuffer` parameter configures the maximum amount of packet buffer, by percentage, that the packets associated with the specified QoS profile can consume. Regardless of the setting for this parameter, the system does not drop any packets as long as packet buffer memory remains available and the current buffer use of the specified QoS profile is below the specified `maxbuffer` setting.

The `weight` parameter configures the relative weighting for each QoS profile. Because each QoS profile has a default weight of 1, all QoS profiles have equal weighting. If you configure a QoS profile with a weight of 4, that specified QoS profile is serviced 4 times as frequently as the remaining QoS profiles, which still have a weight of 1. If you configure all QoS profiles with a weight of 16, each QoS profile is serviced equally but for a longer period.

Example

The following command configures the QoS profile parameters of QoS profile QP1:

```
configure qosprofile qp1 maxbuffer 75 weight 4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure qosscheduler

```
configure qosscheduler [strict-priority | weighted-round-robin]
```

Description

Specifies the method the switch uses to service QoS profiles.

Syntax Description

strict-priority	Specifies the switch services the higher-priority QoS profiles first.
weighted-round-robin	Specifies the switch services all QoS profiles based on the configured weighting for each QoS profile.

Default

Strict-priority.

Usage Guidelines

You configure the QoS scheduling algorithm for the entire switch; you cannot configure QoS scheduling for each port separately.



NOTE

If you specify strict-priority, lower-priority queues are not serviced at all as long as higher-priority queues have any remaining packets. If you specify weighted-round-robin, the switch services higher-weighted queues more frequently but continues to service lower-weighted queues (even when packets remain in the higher-weighted queues).

Example

The following command configures the switch for weighted-round-robin servicing:

```
configure qosscheduler weighted-round-robin
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure traffic ingress queue add egress queue

```
configure traffic ingress queue <queue_name> add egress queue <equeue_name>
ports [all | <port_list>]
```

Description

Using HQoS, associates the egress queue and egress ports to the ingress queue and specifies the egress interfaces. The set of egress ports may be different for every ingress queue to which the egress queue is added.

Syntax Description

queue_name	Enter the ingress traffic queue you are associating with the specified egress queue.
equeue_name	Enter the egress queue you are associating with the specified ingress queue.
all	Specifies all egress ports are associated with the specified association between this ingress queue and this egress queue.
port_list	Enter the egress port numbers you want associated with the specified association between this ingress queue and this egress queue.

Default

N/A.

Usage Guidelines



NOTE

The egress traffic queue is applied to ports only.

You must have configured the HQoS ingress and egress queues prior to using this command to associate the egress queue with an ingress queue. Also, you must have configured the ingress queue to allow egress shaping prior issuing this command. If you issue this command using an ingress-only queue, the command fails, and the system returns the following message:

```
No egress shaping allowed on this ingress queue
```

You associate one egress queue with different ingress queues by re-issuing this command and specifying different ports.

Traffic from the ingress queue that flows to egress ports other than those specified in the egress queue is unlimited on egress.

Example

The following command associates the egress queue tqe with the ingress queue tqe and adds all ports for egress for that association:

```
configure traffic ingress queue tqe add egress queue tqe ports 2:1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

configure traffic ingress queue delete egress queue

```
configure traffic ingress queue <queue_name> delete egress queue [all |
<egress_name>] ports [all | <port_list>]
```

Description

Using HQoS, deletes the egress queue association with the specified ingress queue.

Syntax Description

queue_name	Enter the ingress traffic queue for which you are deleting the association with the specified egress queue(s).
all	Specifies removing the association with all egress queues for the specified ingress queue.
egress_name	Enter the egress queue for which you are deleting the association with the specified ingress queue.
all	Specifies all egress ports are associated with the specified association between this ingress queue and this egress queue.
port_list	Enter the egress port numbers you want associated with the specified association between this ingress queue and this egress queue.

Default

N/A.

Usage Guidelines

Use this command to remove the associated egress queue(s) from the specified ingress queue.

Example

The following command removes all egress queue associations from the ingress queue tq5:

```
configure traffic ingress queue tq5 delete egress queue all
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

configure traffic mode

```
configure traffic mode [pr-and-cr | strict-priority]
```

Description

Configures the switch to use either HQoS bandwidth mode, setting peak and committed rates, or HQoS strict priority mode.



NOTE

You must have an enabled Feature Pack to use the HQoS bandwidth mode.

Syntax Description

pr-and-cr	Specifies that the switch uses bandwidth mode for HQoS. You must have an enabled MEF Feature Pack to use this HQoS mode.
strict-priority	Specifies that the switch uses strict priority mode for HQoS.

Default

Strict priority.

Usage Guidelines

The default HQoS mode is strict priority mode. You must reboot when you choose the bandwidth mode (`configure traffic mode pr-and-cr`) or when you change from one mode to the other. After issuing this command, the system returns the following message:

```
Do you want to reboot?
```

You must also delete all traffic queues prior to changing the HQoS mode.

Example

The following command puts the switch into bandwidth mode for HQoS:

```
configure traffic mode pr-and-cr
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

configure traffic queue aggregate-meter

```
configure traffic queue <queue_name> aggregate-meter <metername>
```

Description

Configures the aggregate rate limit for the HQoS queue by associating the specified queue with the specified meter.

Syntax Description

queue_name	Specifies the traffic queue you are configuring.
metername	Specifies the meter you are associating with this queue, which contains the peak rate for the queue.

Default

Bypass.

Usage Guidelines

Use this command to set the peak rate in strict priority HQoS; if you do not issue this command, the default is to bypass this traffic queue.

You use this command with all types of strict priority HQoS traffic queues: ingress-only, ingress queues that allow egress shaping, and egress queues.

This command creates the aggregate, or peak, rate limit for the specified queue by associating that queue with the meter containing specified rate values.

Example

The following command configures the test traffic queue with an aggregate rate limit specified by the peak rate in the m1_100 meter:

```
configure traffic queue test aggregate-meter m1_100
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch in strict priority mode HQoS.

configure traffic queue cos

```
configure traffic queue <queue_name> cos [{ all<metername>} | {<level>
<metername>} | bypass]
```

Description

In bandwidth mode HQoS, associates a meter to a given CoS value pair sub-queue in bandwidth mode HQoS (to use one meter for all sub-queues in the traffic queue, use the `all` variable).

Syntax Description

queue_name	Specifies the traffic queue you are configuring.
all	Specifies you are associating this meter, which contains the committed and peak rates, with all sub-queues (or CoS value pairs) in the specified traffic queue. Using this variable associates one meter with the entire L2 traffic queue.
metername	Specifies the meter, which contains the committed and peak rates, you are associating either with the entire L2 traffic queue or with the specified L1 sub-queue (or CoS value pair).
level	Specifies the sub-queue, or CoS value pair, you are associating the meter (containing the committed and peak rates) with; choose one of the following four levels, or sub-queues: ● COS0-1 ● COS2-3 ● COS4-5 ● COS6-7
bypass	Specifies to bypass HQoS for the specified sub-queue, or CoS value pair.

Default

N/A.

Usage Guidelines

You use this command with all types of bandwidth mode HQoS traffic queues: ingress-only, ingress queues that allow egress shaping, and egress queues.

A meter should be associated for each of the four CoS value pairs—0,1; 2,3; 4,5; 6,7—and that meter must have configured both committed and peak rates. All four CoS value sub-queues must be configured or bypassed before you can configure any ports for the specified traffic queue. A single meter may be associated with all CoS sub-queues by specifying the `all` keyword and a specific `metername`.

Example

The following command configures all CoS value pairs in the black traffic queue with committed and peak rates specified by the bw1_100 meter:

```
configure traffic queue black cos all bw1_100
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only in bandwidth mode HQoS on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

configure vlan qosprofile

```
configure vlan <vlan_name> {qosprofile} <qosprofile>
```

Description

Configures a VLAN to use a particular QoS profile.

Syntax Description

vlan_name	Specifies a VLAN name.
qosprofile	Specifies a QoS profile.
	NOTE: None is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

Default

The default for the BlackDiamond 10808 and 12800 series switches is QP1, and the default for the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches is none.

Usage Guidelines

Extreme switches support eight QoS profiles (QP1 to QP8) for each port. The QoS profiles QP6 and QP7 are not available to you on a SummitStack.



NOTE

This command applies only to untagged packets on the BlackDiamond 10808 and BlackDiamond 12800 series switches.

Example

The following command configures VLAN *accounting* to use QoS profile QP3:

```
configure vlan accounting qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vman rate-limit flood-traffic designated-port

```
configure vman <vman_name> rate-limit flood-traffic designated-port <port>
```

Description

Sets the vMAN port of which the system uses the rate-limit profile to control the flooding, multicast, and broadcast traffic in a vMAN with egress rate-limiting applied.

Syntax Description

vman_name	Specifies the egress rate-limited enabled vMAN for which you are controlling the flooding, multicast, and broadcast traffic.
port	Specifies the vMAN port of which you are using the rate-limit profile.

Default

N/A.

Usage Guidelines

You enable egress rate limiting for a vMAN by using vMAN ACLs. (See [Chapter 11, “ACL Commands,”](#) for information on vMAN ACLs.)

Once you apply egress rate limiting to a vMAN, you can configure how to control flooding, multicast, and broadcast traffic on the specified vMAN. You specify one port in the vMAN as the designated port, and the system applies the egress rate limiting profile of the designated port to send flood, multicast, and broadcast traffic on all ports on that vMAN.

To display the port you designated to use the rate-limit profile that controls the flooding traffic for the vMAN with egress rate limiting enabled, issue the `show vman detail` command.

The following is sample output from this command (information is displayed on the ERL Designated Port line):

```
VMAN Interface with name vman110 created by user
  Admin State:    Enabled          Tagging:      802.1Q Tag 110
  Priority:       802.1P Priority 0
  Virtual router: VR-Default
  IPv6:          None
  STPD:          None
  Protocol:      Match all unfiltered protocols
  Loopback:      Disabled
  NetLogin:      Disabled
  ERL Designated Port: 6:2
  QosProfile:    QP1
  Ports: 4.      (Number of active ports=4)
    Untag:       *2:1,    *2:4
    Tag:         *6:1,    *6:2
```

Flags: (*) Active, (!) Disabled, (g) Load Sharing port
 (b) Port blocked on the vlan, (a) Authenticated NetLogin port
 (u) Unauthenticated NetLogin port, (m) Mac-Based port

Example

The following command designates port 5:5 as the port you are using the rate-limit profile to control the flooding, multicast, and broadcast traffic in the egress rate-limit enabled vMAN metro:

```
configure vman metro ratelimit flood-traffic designated-port 5:5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

create meter

```
create meter <metername>
```

Description

Creates a meter to limit the rate of packet flow.

Syntax Description

metername	Specifies the meter name.

Default

N/A.

Usage Guidelines

When using meters with ACLs, the meter created with this command is associated with an ACL by specifying the meter name in the action modifier field of the ACL entry. The deny keyword should not be an action for the ACL, or the packet will be dropped regardless of the metering state.

Example

The following command creates the meter *maximum_bandwidth*:

```
create meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches, and the BlackDiamond 12800 R-series switch.

create qosprofile

```
create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Creates a QoS profile.

Syntax Description

QP1....QP7	Specifies the QoS profile you want to create.

Default

N/A.

Usage Guidelines

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches allow dynamic creation and deletion of QoS queues, with Q1 and Q8 always available, rather than the 8 fixed queues on the BlackDiamond 10808 switch.

You must create the QoS profile before you can configure it. or assign it to ports or VLANs. Use the command `configure qosprofile <qosprofile> {maxbuffer <percent>} {weight <value>}` to configure the QoS profile once you create it. Use the `configure ports <port_list> {qosprofile} <qosprofile>` command to associate a QoS profile with ports and the `configure vlan <vlan_name> {qosprofile} <qosprofile>` command to associate a QoS profile with VLANs.

QoS profiles QP1 and QP8 are defaults and cannot be deleted.

Qos profiles QP6 and QP7 can NOT be created in a SummitStack; these two queues are reserved for stacking protocol control traffic.



NOTE

The sFlow application uses QP2 to sample traffic on the BlackDiamond 8800 original modules, SummitStack, and the Summit series switch; any traffic grouping using QP2 may encounter unexpected results when sFlow is enabled on these specific devices.

Example

The following command creates the user-created QoS profile QP3:

```
create qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

create traffic queue allow-egress-shaping

```
create traffic queue <queue_name> allow-egress-shaping {pr-and-cr | strict-  
priority}
```

Description

Creates an ingress traffic queue for use in HQoS rate limiting that allows association with an egress traffic queue. You can use this ingress queue for several egress queues; you configure each association separately. Rate limiting allows you to control bandwidth throughout the network.

Syntax Description

queue_name	Specifies the ingress traffic queue you are configuring that will allow egress shaping.
pr-and-cr	Use optionally when using bandwidth mode HQoS.
queue_name	Use optionally when using strict priority mode HQoS.

Default

N/A.

Usage Guidelines

Use this command for ingress and egress HQoS rate limiting.

Allows the ingress queue that is created with this command to be associated with an egress queue. To accomplish egress rate limiting on this ingress queue, you must create, configure, and associate an egress queue to a specified ingress queue.

You can use this ingress queue for several different egress queues; you must configure each association *separately*.

Look at the counters (using the `show traffic queue [port <port_list> | <queue_name>] statistics` command) to see how the traffic is flowing.

Example

The following command creates the ingress traffic queue that allows egress shaping named tq:

```
create traffic queue tq allow-egress-shaping
```

History

This command was first available in ExtremeXOS 11.4.

The optional `pr-and-cr` was added in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

create traffic queue egress

```
create traffic queue <queue_name> egress {pr-and-cr | strict-priority}
```

Description

Creates an egress traffic queue for use in ingress and egress HQoS rate limiting. To accomplish this egress HQoS, you must have an ingress queue that allows egress shaping to associate with this egress queue. HQoS allows you to control bandwidth throughout the network.

Syntax Description

queue_name	Specifies the egress traffic queue you are configuring for ingress and egress rate limiting.
pr-and-cr	Use optionally when using bandwidth mode HQoS.
queue_name	Use optionally when using strict priority mode HQoS.

Default

N/A.

Usage Guidelines

With ingress and egress HQoS, egress traffic queues are associated with one or more ingress traffic queues. Traffic available for HQoS rate limiting by the egress traffic queue is *only* that traffic was not previously dropped by the associated ingress traffic queue.

To accomplish ingress and egress HQoS rate limiting, you must:

- create and configure an ingress queue that allows egress shaping
- create and configure an egress queue
- explicitly associate an egress queue with an ingress queue (see [create traffic queue allow-egress-shaping](#))

Look at the counters (using the `show traffic queue [port <port_list> | <queue_name>] statistics` command) to see how the traffic is flowing.

Example

The following command creates an egress traffic queue named tqe:

```
create traffic queue tqe egress
```

History

This command was first available in ExtremeXOS 11.4.

The optional `pr-and-cr` was added in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

create traffic queue ingress-only

```
create traffic queue <queue_name> ingress-only {pr-and-cr | strict-
priority}
```

Description

Creates an ingress-only traffic queue for use in ingress-only and ingress and egress HQoS rate limiting. You cannot associate any egress shaping with this traffic queue; you can perform HQoS on ingressing traffic *only*. Rate limiting allows you to control bandwidth throughout the network.

Syntax Description

queue_name	Enter the name of the ingress-only traffic queue you are configuring. This traffic queue rate limits only ingressing traffic.
pr-and-cr	Use optionally when using bandwidth mode HQoS.
queue_name	Use optionally when using strict priority mode HQoS.

Default

N/A.

Usage Guidelines

This command allows you to rate limit the HQoS ingress traffic stream as defined in this queue. You cannot perform any egress HQoS on these ingress-only queues.

Look at the counters (using the `show traffic queue [port <port_list> | <queue_name>] statistics` command) to see how the traffic is flowing.

Example

The following command creates the HQoS ingress-only strict priority traffic queue test:

```
create traffic queue test ingress-only
```

History

This command was first available in ExtremeXOS 11.4.

The optional `pr-and-cr` was added in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

delete meter

```
delete meter <metername>
```

Description

Deletes a meter.

Syntax Description

metername	Specifies the meter name.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the meter *maximum_bandwidth*:

```
delete meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 switches, SummitStack, and the Summit family of switches, and the BlackDiamond 12800 R-series switch.

delete qosprofile

```
delete qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Deletes a user-created QoS profile.

Syntax Description

QP1....QP7	Specifies the user-created QoS profile you want to delete.

Default

N/A.

Usage Guidelines

You cannot delete the default QoS profiles of QP1 and QP8. On a SummitStack, you also cannot delete QoS profiles QP6 and QP7. If you attempt to delete these QoS profiles, the system returns an error.

All configuration information associated with the specified QoS profile is removed.

Example

The following command deletes the user-created QoS profile QP3:

```
delete qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

delete traffic queue

```
delete traffic queue <queue_name>
```

Description

Deletes the specified traffic queue and removes all ingress-only and ingress and egress HQoS rate limiting resources associated with this queue from the hardware. This command does not delete any associated meters, which remain available for use with other traffic queues.

Syntax Description

queue_name	Specifies the traffic queue you are deleting.

Default

N/A.

Usage Guidelines

Prior to deleting an HQoS traffic queue, you must remove all ACL policy file associations; you cannot delete a traffic queue that is currently associated with one or more ACL policy files.

When you delete an ingress traffic queue associated with an egress traffic queue, the associated egress traffic queue remains.

Once you have associated an egress queue to an ingress queue, you cannot delete the egress queue; the system returns an error message. However, prior to associating the egress queue with the ingress queue (see `configure traffic ingress queue add egress queue` command), you can delete the egress queue.

When you delete any type of traffic queue, the associated meters are *not* deleted. Rather, those meters remain and can be associated with other traffic queues. To display the configured meters, issue the `show meters` command.

Example

The following command deletes the test strict priority traffic queue:

```
delete traffic queue test
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

disable diffserv examination port

```
disable diffserv examination port [<port_list> | all]
```

Description

Disables the examination of the DiffServ field in an IP packet.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination should be disabled for all ports.

Default

Disabled.

Usage Guidelines

The diffserv examination feature is disabled by default.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure DiffServ.1p information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.

Example

The following command disables DiffServ examination on selected ports:

```
disable diffserv examination ports 5:3,5:5,6:6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable diffserv replacement port

```
disable diffserv replacement port [<port_list> | all]
```

Description

Disables the replacement of DiffServ code points in packets transmitted by the switch.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply. \
all	Specifies that DiffServ replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The DiffServ replacement feature is disabled by default.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examinational mappings apply on ports in more than one VR. If you attempt to configure DiffServ information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.



NOTE

The specified ports are the ingress ports.

Example

The following command disables DiffServ replacement on selected ports:

```
disable diffserv replacement ports 1:2,5:5,6:6
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable dot1p examination ports

```
disable dot1p examination ports [<port_list> | all]
```

Description

Prevents examination of the 802.1p priority field as part of the QoS configuration.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be disabled for all ports.

Default

Enabled.

Usage Guidelines

The 802.1p examination feature is enabled by default. Once you disable 802.1p examination, the system releases ACL resources; if you re-enable this feature, those ACL resources are once again unavailable. See *ExtremeXOS Concepts Guide* for information on available ACLs on these platforms.

Extreme Networks recommends disabling this feature when you are not using this form of traffic grouping to free resources.

SummitStack only. dot1p examination cannot be disabled for priority values 5 and 6. However, the precedence of the examination is lowered so that all other traffic grouping precedences are higher. The mappings you configure with the `configure dot1p type` command remain in effect.

Example

The following command disables 802.1p value examination on ports 1 to 5:

```
disable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

disable dot1p replacement ports

```
disable dot1p replacement ports [<port_list> | all]
```

Description

Disables the ability to overwrite 802.1p priority values for a given set of ports.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that 802.1p replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the 802.1p feature. If you attempt to configure 802.1p information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.



NOTE

The specified ports are ingress ports.

Example

The following command disables 802.1p value replacement on all ports:

```
disable dot1p replacement ports all
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable diffserv examination port

```
enable diffserv examination port [<port_list> | all]
```

Description

Enables the DiffServ field of an IP packet to be examined in order to select a QoS profile.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination is enabled for all ports.

Default

Disabled.

Usage Guidelines

The diffserv examination feature is disabled by default.

If you are using DiffServ for QoS parameters, Extreme Networks recommends that you also configure 802.1p or port-based QoS parameters to ensure that high-priority traffic is not dropped prior to reaching the MSM on modular switches.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure DiffServ information on a port that is in more than one virtual router, the system returns the following message:

```
Warning: Port belongs to more than one VR. Port properties related to diff serv and
code replacement will not take effect.
```

Example

The following command enables DiffServ examination on selected ports:

```
enable diffserv examination ports 1:1,5:5,6:2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable diffserv replacement ports

```
enable diffserv replacement ports [<port_list> | all]
```

Description

Enables the DiffServ code point to be overwritten in IP packets transmitted by the switch.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The diffserv replacement feature is disabled by default.



NOTE

The port in this command is the ingress port.

Eight user-defined DiffServ code points can be configured on each port. The QoS profile is used to select one of the eight code points.

If you are using DiffServ for QoS parameters, Extreme Networks recommends that you also configure 802.1p or port-based QoS parameters to ensure that high-priority traffic is not dropped prior to reaching the MSM on modular switches.



NOTE

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure DiffServ information on a port that is in more than one virtual router, the system returns the following message:

```
Warning: Port belongs to more than one VR. Port properties related to diff serv and
code replacement will not take effect.
```

Example

The following command enables DiffServ replacement on selected ports:

```
enable diffserv replacement ports 5:3,5:5,6:2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable dot1p examination ports

```
enable dot1p examination ports [<port_list> | all]
```

Description

Allows the 802.1p priority field to be examined as part of the QoS configuration.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
all	Specifies that dot1p examination should be enabled for all ports.

Default

Enabled.

Usage Guidelines

To increase available ACLs, you can disable the 802.1p examination feature if you are not running QoS or are running QoS using DiffServ. See *ExtremeXOS Concepts Guide* for information on ACL limitations on these platforms.

Use this command to re-enable the 802.1p examination feature.

Example

The following command enables dot1p examination on ports 1 to 5:

```
enable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

enable dot1p replacement ports

```
enable dot1p replacement ports [<port_list> | all]
```

Description

Allows the 802.1p priority field to be overwritten on egress according to the QoS profile to 802.1p priority mapping for a given set of ports.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet.



NOTE

The port in this command is the ingress port.

If 802.1p replacement is enabled, the 802.1p priority information that is transmitted is determined by the hardware queue that is used when transmitting the packet.



NOTE

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the 802.1p feature. If you attempt to configure 802.1p information on a port that is in more than one virtual router, the system returns the following message:

```
Warning: Port belongs to more than one VR. Port properties related to diff serv and
code replacement will not take effect.
```

Example

The following command enables dot1p replacement on all ports:

```
enable dot1p replacement ports all
```

History

This command was available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show access-list meter

```
show access-list meter {<metername>} [any | ports <portlist> | vlan
<vlanname>]
```

Description

Displays the specified access list meter statistics and configurations.

Syntax Description

metername	Specifies the ACL meter to display.
portlist	Specifies to display the meters on these ports.
vlanname	Specifies to display the meters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the ACL meters.

Example

The following example displays all the counters the ACL on port 7:1

```
show access-list meter port 2:1
```

The output of this command is similar to the following:

ACL Name	Vlan	Committed Rate (Kbps)	Port Max Burst Size (K)	Out-of-Profile Action	Out-of-Profile DSCP	Out-of-Profile Packet Count
donacl	*		7:1			
m1		10000	Max	Drop		0

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

show diffserv

```
show diffserv [examination | replacement]
```

Description

Displays the DiffServ-to-QoS profile mapping.

Syntax Description

examination	Specifies to display currently configured mapping for each DiffServ code point to QoS profile.
replacement	Specifies to display currently configured code point replacement value.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the "->" in the display (shown below) becomes "*>".

Examples

The following command displays the current QoS-to-DiffServ mappings for each code point on the switch:

```
show diffserv examination
```

The following command displays the current DiffServ-to-QoS mappings on the switch:

```
show diffserv replacement
```

Because the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches have 2 default QoS profiles and the BlackDiamond 10808 and 12800 series switches has 8 default QoS profiles you see different displays depending on the platform.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Following is sample output from the `show diffserv examination` command on the BlackDiamond 8810 switch:

CodePoint->QOSProfile mapping:

```
00->QP1 01->QP1 02->QP1 03->QP1 04->QP1 05->QP1 06->QP1 07->QP1
08->QP1 09->QP1 10->QP1 11->QP1 12->QP1 13->QP1 14->QP1 15->QP1
16->QP1 17->QP1 18->QP1 19->QP1 20->QP1 21->QP1 22->QP1 23->QP1
24->QP1 25->QP1 26->QP1 27->QP1 28->QP1 29->QP1 30->QP1 31->QP1
32->QP1 33->QP1 34->QP1 35->QP1 36->QP1 37->QP1 38->QP1 39->QP1
40->QP1 41->QP1 42->QP1 43->QP1 44->QP1 45->QP1 46->QP1 47->QP1
48->QP1 49->QP1 50->QP1 51->QP1 52->QP1 53->QP1 54->QP1 55->QP1
56->QP8 57->QP8 58->QP8 59->QP8 60->QP8 61->QP8 62->QP8 63->QP8
```

Following is sample output from the `show diffserv replacement` command on the BlackDiamond 8810 switch:

```
QOSProfile->CodePoint mapping:
  QP1->00
  QP8->56
```

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure DiffServ information on a port that is in more than one virtual router, the system returns the following message:

Warning: Port belongs to more than one VR. Port properties related to diff serv and code replacement will not take effect.

Following is sample output from the `show diffserv examination` command on the BlackDiamond 10808 switch:

```
CodePoint->QOSProfile mapping:
  00->QP1 01->QP1 02->QP1 03->QP1 04->QP1 05->QP1 06->QP1 07->QP1
  08->QP2 09->QP2 10->QP2 11->QP2 12->QP2 13->QP2 14->QP2 15->QP2
  16->QP3 17->QP3 18->QP3 19->QP3 20->QP3 21->QP3 22->QP3 23->QP3
  24->QP4 25->QP4 26->QP4 27->QP4 28->QP4 29->QP4 30->QP4 31->QP4
  32->QP5 33->QP5 34->QP5 35->QP5 36->QP5 37->QP5 38->QP5 39->QP5
  40->QP6 41->QP6 42->QP6 43->QP6 44->QP6 45->QP6 46->QP6 47->QP6
  48->QP7 49->QP7 50->QP7 51->QP7 52->QP7 53->QP7 54->QP7 55->QP7
  56->QP8 57->QP8 58->QP8 59->QP8 60->QP8 61->QP8 62->QP8 63->QP8
```

Following is sample output from the `show diffserv replacement` command on the BlackDiamond 10808 switch:

```
QOSProfile->CodePoint mapping:
  QP1->00
  QP2->08
  QP3->16
  QP4->24
  QP5->32
  QP6->40
  QP7->48
  QP8->56
```

History

This command was first available in ExtremeXOS 10.1.

The variables were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show dot1p

```
show dot1p
```

Description

Displays the 802.1p-to-QoS profile mappings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the current 802.1p-to-QoS mappings on the switch:

```
show dot1p
```

Because the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches have 2 default QoS profiles and the BlackDiamond 10808 and 12800 series switches has 8 default QoS profiles you see different displays depending on the platform.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. Following is sample output from the `show dot1p` command on the BlackDiamond 8810 switch:

802.1p	Priority	Value	QOS	Profile
	0		QP1	
	1		QP1	
	2		QP1	
	3		QP1	
	4		QP1	
	5		QP1	
	6		QP1	
	7		QP8	

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the 802.1p feature. If you attempt to configure 802.1p information on a port that is in more than one virtual router, the system returns the following message:

```
Warning: Port belongs to more than one VR. Port properties related to diff serv and
code replacement will not take effect.
```

Following is sample output from the `show dot1p` command on the BlackDiamond 10808 switch:

802.1p	Priority	Value	QOS	Profile
	0		QP1	
	1		QP2	
	2		QP3	
	3		QP4	
	4		QP5	
	5		QP6	
	6		QP7	
	7		QP8	

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show meter

```
show meter
```

Description

Displays the configured meters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays meters on the switch:

```
show meter
```

Display on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only.



NOTE

When you are using a BlackDiamond 8800 series switches, SummitStack, or the Summit family of switches, you configure a peak rate for QoS meters using the `configure meter <metername> {max-burst-size <burst-size> [Gb | Kb | Mb]} {committed-rate <cir> [Gbps | Mbps | Kbps]} {out-actions [drop | set-drop-precedence {dscp [none | <dscp-value>}]}` command.

The following is sample output from this command:

```
-----
Name      Committed Rate(Kbps)  Peak Rate(Kbps)
-----
peggy          1000000          --
```

Display on the BlackDiamond 12800 R-series only.



NOTE

When you are using a BlackDiamond 12800 R-series switch, you configure a peak rate for strict priority HQoS meters using the `configure meter <metername> [peak-rate <pr> [Gbps | Mbps | Kbps]]` and a committed and peak

rate for bandwidth mode HQoS meters using the `configure meter <metername> [committed-rate <cir> [Gbps | Mbps | Kbps] | peak-rate <pr>[Gbps | Mbps | Kbps]]` command.

The following is sample output from this command for strict priority mode HQoS on the BlackDiamond 12800 R-series switch:

```
BD-12804.2 # show meter
```

Name	Committed Rate(Kbps)	Peak Rate(Kbps)
dummy	--	10000000
m1_1	--	1000
m1_10	--	10000
m1_100	--	100000
m1_11	--	11000
m1_12	--	12000
m1_13	--	13000
m1_14	--	14000
m1_145	--	145000
m1_15	--	15000
m1_150	--	150000
m1_16	--	16000
m1_17	--	17000
m1_18	--	18000
m1_19	--	19000
m1_2	--	2000
m1_20	--	20000

The following is sample output from this command for strict priority mode HQoS on the BlackDiamond 12800 R-series switch:

```
BD-12804.26 # show meter
```

Name	Committed Rate(Kbps)	Peak Rate(Kbps)	Burst Size(KB)
Gig1	1000000	1000000	--
m100	100000	150000	--
m150	150000	200000	--
m200	200000	250000	--
m250	0	250000	--
m50	50000	100000	--
m500	0	500000	--
m750	0	750000	--
m950	0	950000	--
me250	150000	250000	--
me500	300000	500000	--
me750	500000	750000	--
me950	950000	950000	--
Meg100	100000	100000	--
Meg200	200000	200000	--
Meg250	250000	250000	--

History

This command was first available in ExtremeXOS 11.4.

Bandwidth mode HQoS was added in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches, SummitStack, the Summit family of switches, and the BlackDiamond 12800 R-series switch (an enabled MEF Feature Pack is required for bandwidth mode HQoS on the BlackDiamond 12800 R-series switch).

show ports qosmonitor

```
show ports <port_list> qosmonitor {ingress | egress} {bytes | packets} {no-
refresh}
```

Description

Displays real-time (or snapshot) ingress rate or shaping statistics in addition to QoS statistics for egress packets on one or more ports.

Syntax Description

port_list	Specifies one or more slots and ports.
egress	Specifies to display statistics in egress. Default.
ingress	Specifies to display statistics on ingress rate shaping.
bytes	Specifies to display statistics by bytes.
packets	Specifies to display statistics by packets.
no-refresh	Specifies a static snapshot of data.

Default

Shows QoS statistics for egress unless you specify ingress; shows real-time or snapshot statistics.

Usage Guidelines

The real-time snapshot scrolls through the given port list to provide statistics. You toggle, using the spacebar, to show statistics by packets or bytes.

If you do not specify a port number or range of ports, collision statistics are displayed for all ports. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command. (The toggle functionality does not work with the `no-refresh` variable, which provides a snapshot of statistics; you must use the variable `bytes` to display the bytes in snapshot.)

Examples

The following command shows the QoS egress statistics related to the specified ports:

```
# show port 1:1-1:2 qosmonitor
```

Following is sample output from this command:

```
Qos Monitor Req Summary                               Thu Mar  2 10:58:23 2006
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
          Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
          Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts
=====
1:1         0         0         0         0         0         0         0         0
1:2         0         0         0         0         0         0         0         0
=====
Spacebar->Toggle screen 0->Clear counters  U->Page up   D->Page down ESC->exit
```

The following command shows the QoS ingress statistics related to the specified ports:

```
# show port 1:1-1:2 qosmonitor ingress
```

Following is sample output from this command:

```
Qos Monitor Req Summary                               Thu Mar  2 10:59:28 2006
Port      IQP1      IQP2      IQP3      IQP4      IQP5      IQP6      IQP7      IQP8
          Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
          Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts
=====
1:1         0         0         0         0         0         0         0         0
1:2         0         0         0         0         0         0         0         0
=====
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0

Also, you must specify the ports in ExtremeXOS 11.0.

The egress and no-refresh keywords were added in ExtremeXOS 11.3.

The bytes and packets keywords, as well as the toggling functionality, were added in ExtremeXOS 11.4.

Platform Availability

This command is available on the BlackDiamond 10808, 12800 series, SummitStack, and Summit X450e switches.

show qosprofile

```
show qosprofile [ingress | egress | ports [ all | <port_list>]
```

Description

Displays QoS information on the switch.

Syntax Description

ingress	Specifies ingress queues. NOTE: This parameter is available only on the BlackDiamond 10808 switch.
egress	Specifies egress queues; this is the default value. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches.
ports	Specifies to display information for specified ports. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond a-series and e-series modules, and the Summit X450a and X450e series switches.
port_list	Specifies a list of slots and ports. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond a-series and e-series modules, and the Summit X450a and X450e series switches.
all	Specifies all ports. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond a-series and e-series modules, and the Summit X450a and X450e series switches.

Default

Displays egress QoS information for all ports.

Usage Guidelines

The displayed QoS profile information differs depending on the platform you are running on.

BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches only, whether or not connected in a SummitStack. The following information displays for these switches:

- Available Qos profiles
- Weight
- Percentage of maximum buffer

When you specify ports, the following information also displays per QoS queue:

- Committed rate
- Peak rate

- Minimum bandwidth
- Maximum bandwidth

BlackDiamond 8800 original modules and Summit X450 series switch only, whether or not connected in a SummitStack. The following information displays for these switches:

- Available Qos profiles
- Weight
- Percentage of maximum buffer

BlackDiamond 10808K and 12800 series switches only. If you do not specify the variable `ingress`, the system displays the information for egress QoS values on the BlackDiamond 10808 switch (ingress is not supported on the BlackDiamond 12800 series switches).

The following information displays for the BlackDiamond 10808 and 12800 series switches for each QoS profile:

- Qos profile
- Minimum bandwidth
- Maximum bandwidth
- Priority

Example

The display varies depending on your platform.

BlackDiamond 8800 a-series and e-series modules and Summit X450a and X450e series switches only, whether or not included in a SummitStack. The following sample output displays for these switches when you omit the optional `port` parameter:

```
QP1    Weight = 1    Max Buffer Percent = 100
QP8    Weight = 1    Max Buffer Percent = 100
```

When you add the optional `port` parameter, the switch displays the following sample output:

```
Port: 4:1
QP1    CIR    =      10M PR      =    100M
QP8    MinBw  =      0% MaxBw  =    100%
```



NOTE

This last sample output is not available on the XGS2 ports.

BlackDiamond 8800 original modules and Summit X450 series switch only, whether or not connected in a SummitStack. No arguments are available on the BlackDiamond 8800 original modules or the Summit X450 series switch for this command. Following is sample output from the `show qosprofile` command:

```
QP1    Weight = 1    Max Buffer Percent = 100
QP2    Weight = 4    Max Buffer Percent = 90
QP8    Weight = 1    Max Buffer Percent = 100
```

BlackDiamond 10808 and 12800 series switches only. You can display the egress QoS profiles on the BlackDiamond 10808 and 12800 series switches. The following shows sample output for the `show qosprofile ports 1:1` command, which displays the egress QoS profiles:

```
Port: 1:1
  QP1  MinBw=  0% MaxBw=100% Pri=1
  QP2  MinBw=  0% MaxBw=100% Pri=2
  QP3  MinBw=  0% MaxBw=100% Pri=3
  QP4  MinBw=  0% MaxBw=100% Pri=4
  QP5  MinBw=  0% MaxBw=100% Pri=5
  QP6  MinBw=  0% MaxBw=100% Pri=6
  QP7  MinBw=  0% MaxBw=100% Pri=7
  QP8  MinBw=  0% MaxBw=100% Pri=8
```

BlackDiamond 10808 switch only. The number of available ingress QoS profiles on the BlackDiamond 10808 switch differs depending on the module you are displaying. Ingress queues are available only on the BlackDiamond 10808 switch; not on the BlackDiamond 12800 series switches.

The following shows sample output for the `show qosprofile ingress ports 8:1` command, which displays the ingress QoS profiles on a BlackDiamond 10808 10 G port:

```
Port: 8:1
  IQP1 MinBw=  0% MaxBw=100% Pri=1
  IQP2 MinBw=  0% MaxBw=100% Pri=2
  IQP3 MinBw=  0% MaxBw=100% Pri=3
  IQP4 MinBw=  0% MaxBw=100% Pri=4
  IQP5 MinBw=  0% MaxBw=100% Pri=5
  IQP6 MinBw=  0% MaxBw=100% Pri=6
  IQP7 MinBw=  0% MaxBw=100% Pri=7
  IQP8 MinBw=  0% MaxBw=100% Pri=8
```

The following shows sample output for the `show qosprofile ingress ports 2:1` command, which displays the ingress QoS profiles on a BlackDiamond 10808 1G port:

```
Port: 2:1
  IQP1 MinBw=  0% MaxBw=100% Pri=1
  IQP2 MinBw=  0% MaxBw=100% Pri=2
```

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

The parameter `ingress` is available only on the BlackDiamond 10808 switch. The parameter `egress` is available only on the BlackDiamond 10808 and 12800 series switches.

The parameter `ports` is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond 8800 a-series and e-series modules, and the Summit X450a and X450e series switches, whether or not included in a SummitStack.

show traffic mode

```
show traffic mode
```

Description

Use this command to view the HQoS traffic mode configured on your switch.

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

Allows you to display the HQoS traffic mode configuration on your switch.

Example

The following command displays the HQoS traffic mode configuration for your switch:

```
show traffic mode
```

The following is an example of the output of this command:

```
BD-12804.28 # show traffic mode
Rate Limiting Mode: PR and CR
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

show traffic queue

```
show traffic queue {detail | <queue_name>}
```

Description

Use this command to view the current ingress-only and ingress and egress HQoS configuration on your switch.

Syntax Description

detail	Specifies displaying greater detail about all configured traffic queues.
queue_name	Specifies displaying greater detail about the specified traffic queue.

Default

N/A.

Usage Guidelines

Allows you to display the HQoS configuration on your switch, as follows:

- Names of traffic queues
- Direction of traffic queues
- Allow egress shaping
- HQoS mode
- Associated queue meter

Example

The following command displays the rate limiting configuration for your switch:

```
show traffic queue
```

The following is an example of the output of this command in strict priority mode HQoS:

```
-----  
Traffic Queue   Direction      Shape Egress  
-----  
limit           Egress         N/A  
peggy           Ingress        Yes  
sample          Ingress        No
```

The following is an example of the output of this command in bandwidth mode HQoS:

```
*BD-12804.27 # show traffic queue
```

Traffic Queue	Direction	Shape Egress	Mode
eq1	Egress	N/A	pr-and-cr
eq2	Egress	N/A	pr-and-cr
eq3	Egress	N/A	pr-and-cr
eq4	Egress	N/A	pr-and-cr
eq5	Egress	N/A	pr-and-cr
iq1	Ingress	Yes	pr-and-cr
iq2	Ingress	Yes	pr-and-cr
iq3	Ingress	Yes	pr-and-cr
iq4	Ingress	Yes	pr-and-cr

The following command displays more detail about each HQoS traffic queue:

```
show traffic queue detail
```

The following is an example of the output of this command:

```
Traffic Queue: limit
  Direction:      Egress
  Mode:           strict-priority
  Queue Meter:    m3_300
Traffic Queue: peggy
  Direction:      Ingress
  Shape Egress:   No
  Mode:           strict-priority
  Queue Meter:    m1_100
Traffic Queue: sample
  Direction:      Ingress
  Shape Egress:   Yes
  Mode:           strict-priority
  Queue Meter:    m2_200
Traffic Queue: iqq2
  Direction:      Ingress
  Shape Egress:   Yes
  Egress Queue:   eqq1      Egress Ports: 1:2
  Mode:           strict-priority
  Queue Meter:    m3_300
```

To display the detail on a single HQoS traffic queue, use the following command:

```
show traffic queue <queue_name>
```

History

This command was first available in ExtremeXOS 11.4.

Bandwidth mode was added in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

show traffic queue statistics

```
show traffic queue [port <port_list> | <queue_name>] statistics
```

Description

Displays packets passed and dropped in high, medium, and low portion of each HQoS traffic queue and by port, if desired.

Syntax Description

port_list	Enter the ports for which you want to display the rate limiting statistics.
queue_name	Enter the name of the traffic queue for which you want to display the rate limiting statistics.

Default

N/A.

Usage Guidelines

Allows you to view the number of packets dropped and passed in each part of the HQoS strict priority queues: high, medium, or low.

If you issue the command, `configure ports rate-limit packet byte-adjustment`, the system calculates uses that byte size to calculate the statistics. If you use this command to decrease by 4, the output of the `show traffic queue statistics` command displays only bytes; the packets are displayed as 0.

Displays the following information:

- Name of traffic queue
- Direction of queue
- Priority (high, medium, low)
- Packets passed
- Bytes passed
- Packets dropped
- Bytes dropped

You can also display this information by port.

Example

The following command displays the rate limiting statistics for the specified port:

```
show traffic queue port 6:2 statistics
```

The following is an example of the output of this command in strict priority mode HQoS:

Traffic Queue Name	D	COS	Passed		Dropped	
			Packets	Bytes	Packets	Bytes
Cust_01	I	High	1492134	131307792	737901	72935288
	I	Med	903345	82316776	43198454	7639238004
	I	Low	7298	863324	1491673790	12672935288
	I	*	2402777	214487892	1535609145	20385108580
po_62_egr	E	High	3472754	331359794	0	0
	E	Med	904283	78810756	0	0
	E	Low	7549233	452274980	3790	135288
	E	*	11926270	8622445530	3790	135288

(*)-PR aggregate CoS Level for entire traffic queue, (I)-Ingress, (E)-Egress

The following is an example of the output of this command in bandwidth mode HQoS:

```
*BD-12804.1 # show traffic queue eq2 statistics
```

Traffic Queue Name	COS		Passed		Dropped	
	Level		Packets	Bytes	Packets	Bytes
eq2	0-1	C	0	0	0	0
		P	0	0	0	0
	2-3	C	0	0	0	0
		P	0	0	0	0
	4-5	C	0	0	0	0
		P	0	0	0	0
	6-7	C	0	0	0	0
		P	0	0	0	0

The following command displays the rate limiting statistics for the specified traffic queue:

```
show traffic queue tq1 statistics
```

The following is an example of the output of this command:

Traffic Queue Name	D	COS	Passed		Dropped	
			Packets	Bytes	Packets	Bytes
tq1	I	High	0	0	0	0
	I	Med	0	0	0	0
	I	Low	0	0	0	0
	I	*	0	0	0	0

(*)-PR aggregate CoS Level for entire traffic queue, (I)-Ingress, (E)-Egress

History

This command was first available in ExtremeXOS 11.4.

Bandwidth mode was added in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

show traffic queue utilization

```
show traffic queue [port <port_list> | <queue_name>] utilization
```

Description

Displays percentage of configured HQoS aggregated traffic each priority level—high, medium, or low—of the specified traffic queue (or port) is using. This value is a percentage of the traffic queue's configured peak rate in each of the priority levels.

Syntax Description

port_list	Enter the ports for which you want to display how much each level is using.
queue_name	Enter the name of the traffic queue for which you want to display how much each level is using.

Default

N/A.

Usage Guidelines

Allows you to view the percentage of the configured HQoS aggregated peak rate traffic in each part of the strict priority queues: high, medium, or low. The utilization percentage is averaged over a 5-second interval.

If you issue the command, `configure ports rate-limit packet byte-adjustment`, the system calculates uses that byte size to calculate the utilization values.

Displays the following information:

- Name of traffic queue
- Direction of queue
- Priority (high, medium, low)
- Percentage of configured peak rate in each of the priority levels

You can also display this information by port.

Example

The following command displays HQoS utilization for the specified port:

```
show traffic queue port 2:1 utilization
```

The following is an example of the output of this command for strict priority mode HQoS:

Traffic		Priority Levels			
Queue Name	D	H	M	L	
-----	-	---	---	---	
tq2	I	22	15	44	
tq3	I	5	7	238	

I=Ingress, E=Egress

The following is an example of the output of this command for bandwidth mode HQoS:

```
* BD-12804.2 # show traffic queue eq2 utilization
```

Traffic		COS (% of configured CIR)				
Queue Name	D	0/1	2/3	4/5	6/7	

eq2	E	--	--	--	--	

I=Ingress, E=Egress

The following command displays HQoS utilization for the specified traffic queue:

```
show traffic queue tq2 utilization
```

The following is an example of the output of this command:

Traffic		Priority Levels			
Queue Name	D	H	M	L	
-----	-	---	---	---	
tq2	I	0	0	105	

I=Ingress, E=Egress

History

This command was first available in ExtremeXOS 11.4.

Bandwidth mode was added in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch; bandwidth mode is available only with an enabled MEF Feature Pack.

unconfigure diffserv

```
unconfigure diffserv [examination | replacement]
```

Description

Uses the default DiffServ examination code point or uses the default DiffServ replacement mapping.

Syntax Description

examination	Specifies to unconfigure the DiffServ examination point.
replacement	Specifies to unconfigure the DiffServ replacement mapping.

Default

N/A.

Usage Guidelines

BlackDiamond 10808 and 12800 series switches only. If a port is in more than one virtual router, you cannot use the DiffServ feature. The default VLAN DiffServ examination mappings apply on ports in more than one VR. If you attempt to configure DiffServ information on a port that is in more than one virtual router, the system returns the following message:

```
Warning: Port belongs to more than one VR. Port properties related to diff serv and  
code replacement will not take effect.
```

Example

The following command removes DiffServ code point examination from ports 5:5-5:8 on a modular switch:

```
unconfigure diffserv examination
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure ports aggregate-bandwidth

```
unconfigure ports <port_list> aggregate-bandwidth
```

Description

Use this command to delete the port-based egress rate limiting parameters; this command returns the port to 100%, or line rate.

Syntax Description

port_list	Specifies the ports from which you want to delete port-based egress rate limiting.

Default

N/A.

Usage Guidelines

Use this command to delete port-based egress rate limiting on the specified ports. This command returns the specified ports to run at 100% or line rate.

Example

The following command deletes port-based egress rate limiting from ports 1:1 through 1:5 on your switch:

```
unconfigure ports 1:1-5 aggregate-bandwidth
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

unconfigure qosprofile

```
unconfigure qosprofile {ingress | egress} {ports [<port_list>|all]}
```

Description

Returns the ingress rate shaping parameters, which is an ingress QoS profile, or the egressing QoS values to default values.

Syntax Description

ingress	Specifies all ingress QoS profiles for specified ports. The default ingress values are as follows: • for 1G I/O modules—IQP1 and IQP2 • for 10G I/O modules—IQP1 to IQP8 If you do not specify ingress or egress, the command returns all egress QoS values to default. NOTE: This parameter is available only on the BlackDiamond 10808 switch.
egress	Specifies an egress QoS profile for specified ports. The default egress values are as follows: • QP1—1, Low (the lowest priority) • QP2—2, LowHi • QP3—3, Normal • QP4—4, NormalHi • QP5—5, Medium • QP6—6, MediumHi • QP7—7, High • QP8—8, HighHi (highest priority) If you do not specify ingress or egress, the command returns all <i>egress</i> QoS values to default. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches.
ports	Specifies unconfiguring the QoS profiles on specified ports. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond 8800 a-series and e-series modules, and the Summit X450a and X450e series switches, whether or not included in a SummitStack.
port_list	Specifies a list of slots and ports to which the parameters apply. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond 8800 a-series and e-series modules, and the Summit X450a and X450e series switches, whether or not included in a SummitStack.
all	Specifies that this applies to all ports on the device. NOTE: This parameter is available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond 8800 a-series and e-series modules, and the Summit X450a and X450e series switches, whether or not included in a SummitStack.

Default

Default values differ by platform.

BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches only. The default values for the two default egressing QoS profiles (QP1 and QP8) on the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches are shown in this section; the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support ingressing QoS profiles.

- Weight—1
- Maximum buffer—100%

For the BlackDiamond 8800 a-series and e-series modules and the Summit X450a and X450e series switches, whether or not included in a SummitStack, this command also resets the minimum bandwidth to 0% and the maximum bandwidth to 100% for the specified egressing

BlackDiamond 10808 and 12800 series switches only. the default values for egressing QoS profiles on the BlackDiamond 10808 and 12800 series switches are shown in this section:

- QP1 - 1, Low
- QP2 - 2, LowHi
- QP3 - 3, Normal
- QP4 - 4, NormalHi
- QP5 - 5, Medium
- QP6 - 6, MediumHi
- QP7 - 7, High
- QP8 - 8, HighHi (highest priority)

BlackDiamond 10808 switch only. the default values for ingressing QoS profiles on the BlackDiamond 10808 switch are shown in this section.

Ingress QoS is disabled by default.

- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Minimum committed and peak rates—1 Kbps
- Maximum committed and peak rates—varies by I/O module:
 - 1G I/O module—1,000 Mbps
 - 10G I/O module—10,000 Mbps
- Ingress priority—By default, each QoS profile is assigned a different priority level, which varies by I/O module:
 - 1G I/O module:
 - IQP1 - 1 Low
 - IQP2 - 1, LowHi
 - 10G I/O module:
 - IQP1 - 1, Low
 - IQP2 - 2, LowHi
 - IQP3 - 3, Normal

- IQP4 - 4, NormalHi
- IQP5 - 5, Medium
- IQP6 - 6, MediumHi
- IQP7 - 7, High
- IQP8 - 8, HighHi (highest priority)

Usage Guidelines

BlackDiamond 8800 a-series and e-series modules and Summit X450a, X450e, and X250e switches only, whether or not included in a SummitStack. This command also resets the minimum bandwidth to 0% and the maximum bandwidth to 100% for egressing QoS profiles on the specified ports. There are no ingress queues on the BlackDiamond 8800 series switches, SummitStack, and the Summit X450 series switches.

BlackDiamond 10808 switch only. The number of ingress queues per port varies between the 1 Gbps I/O module and the 10 Gbps module. On the 1 Gbps module, you have two ingress queues per port. The priority values of 1 to 4 map to the first queue, and the priority values of 5 to 8 map to the second queue. On the 10 Gbps module, you have eight ingress queues per port. The priority values of 1 to 8 map one to each of the eight queues.

Example

The following command resets the QoS profiles for all ports to default settings:

```
unconfigure qosprofile
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

The ingress parameter is available only on the BlackDiamond 10808 switch.

The parameters egress and ports are available only on the BlackDiamond 10808 and 12800 series switches, the BlackDiamond a-series and e-series modules, and the Summit X450a, X450e, and X250e series switches, whether or not included in a SummitStack.

unconfigure vman rate-limit flood-traffic designated-port

```
unconfigure vman <vman_name> rate-limit flood-traffic designated-port
```

Description

Removes the control on flooding, multicast, and broadcast traffic in a vMAN with egress rate-limiting applied.

Syntax Description

vman_name	Specifies the egress rate-limited enabled vMAN for which you are removing the control on flooding, multicast, and broadcast traffic.

Default

N/A.

Usage Guidelines

Use this command to remove the control on flooding, multicast, and broadcast traffic in the vMAN with egress rate limiting.

To display the port you designated to use the rate-limit profile that controls the flooding traffic for the vMAN with egress rate limiting enabled, issue the `show vman detail` command.

The following is sample output from this command (information is displayed on the ERL Designated Port line):

```
VMAN Interface with name vman400 created by user
  Admin State:      Enabled      Tagging:      802.1Q Tag 400
  Priority:          802.1P Priority 0
  Virtual router:   VR-Default
  IPv6:             None
  STPD:             None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disabled
  NetLogin:         Disabled
  ERL Designated Port: Not configured
  QosProfile:       QP1
  Ports:           0.           (Number of active ports=0)
```

Example

The following command removes the control of flooding, multicast, and broadcast traffic in the egress rate-limit enabled vMAN metro:

```
unconfigure vman metro ratelimit flood-traffic designated-port
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the BlackDiamond 12800 R-series switch.

19 Security Commands

This chapter describes commands for:

- Managing the switch using SSH2
- Configuring switch user authentication through a RADIUS client
- Configuring switch user authentication through TACACS+
- Protecting the switch from Denial of Service attacks

SSH

Secure Shell 2 (SSH2) is a feature of ExtremeXOS that allows you to encrypt session data between a network administrator using SSH2 client software and the switch. Configuration and policy files may also be transferred to the switch using the Secure Copy Program 2 (SCP2).

User Authentication

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administering access to network nodes. The ExtremeXOS RADIUS client implementation allows authentication for SSH2, Telnet or console access to the switch.

Extreme switches are also capable of sending RADIUS accounting information. You can configure RADIUS accounting servers to be the same as the authentication servers, but this is not required.

Terminal Access Controller Access Control System Plus (TACACS+) is a mechanism for providing authentication, authorization, and accounting on a centralized server, similar in function to the RADIUS client. The ExtremeXOS version of TACACS+ is used to authenticate prospective users who are attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.



NOTE

You cannot use RADIUS and TACACS+ at the same time.

Denial of Service

You can configure ExtremeXOS to protect your Extreme switches in the event of a denial of service attack. During a typical denial of service attack, the CPU on the switch gets flooded with packets from multiple attackers, potentially causing the switch to fail. To protect against this type of attack, you can configure the software so that when the number of packets received is more than the configured threshold limit of packets per second, a hardware ACL is enabled.

clear ip-security anomaly-protection notify cache

```
clear ip-security anomaly-protection notify cache {slot [<slot> | all ]}
```

Description

Clear the local protocol anomaly event cache.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

N/A.

Usage Guidelines

This command clears the local protocol anomaly event cache.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

clear ip-security dhcp-snooping entries

```
clear ip-security dhcp-snooping entries { vlan } <vlan_name>
```

Description

Clears the DHCP binding entries present on a VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server.

Default

N/A.

Usage Guidelines

Use this command to clear the DHCP binding entries present on a VLAN. When an entry is deleted, all its associated entries (such as source IP lockdown, secured ARP, and so on) and their associated ACLs, if any, are also deleted.

Example

The following command clears the DHCP binding entry temporary from the VLAN:

```
clear ip-security dhcp-snooping entries temporary
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

clear ip-security source-ip-lockdown entries ports

```
clear ip-security source-ip-lockdown entries ports [ <ports> | all ]
```

Description

Clears locked-down source IP addresses on a per-port basis.

Syntax Description

ports	Specifies the port or ports to be cleared.
all	Specifies that all ports are to be cleared.

Default

N/A.

Usage Guidelines

Use this command to clear locked-down source IP addresses on a per port basis. This command deletes the entries on the indicated ports and clears the associated ACLs.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

clear vlan dhcp-address-allocation

```
clear vlan <vlan_name> dhcp-address-allocation [[all {offered | assigned |
declined | expired}] | <ipaddress>]
```

Description

Removes addresses from the DHCP allocation table.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server.
all	Specifies all IP addresses, or all IP addresses in a particular state.
offered	Specifies IP addresses offered to clients.
assigned	Specifies IP addresses offered to and accepted by clients.
declined	Specifies IP addresses declined by clients
expired	Specifies IP addresses whose lease has expired and not renewed by the DHCP server.
ipaddress	Specifies a particular IP address.

Default

N/A.

Usage Guidelines

You can either delete a single entry, using the IP address, or all entries. If you use the `all` option, you can additionally delete entries in a specific state.

Example

The following command removes all the declined IP addresses by hosts on the VLAN *temporary*:

```
clear vlan temporary dhcp-address-allocation all declined
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure dos-protect acl-expire

```
configure dos-protect acl-expire <seconds>
```

Description

Configures the denial of service protection ACL expiration time.

Syntax Description

seconds	Specifies how long the ACL is in place.

Default

The default is 5 seconds.

Usage Guidelines

This command configures how long the DoS protection ACL remains in place.

Example

This example sets the ACL expiration time to 15 seconds:

```
configure dos-protect acl-expire 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect interval

```
configure dos-protect interval <seconds>
```

Description

Configures the denial of service protection interval.

Syntax Description

seconds	Specifies how often the DoS protection counter is monitored.

Default

The default is one second.

Usage Guidelines

This command configures how often the DoS protection counter is monitored.

Example

This example sets the interval to 5 seconds:

```
configure dos-protect interval 5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect trusted ports

```
configure dos-protect trusted-ports
[ports [<ports> | all]
| add-ports [<ports-to-add> | all]
| delete-ports [<ports-to-delete> | all]
]
```

Description

Configures the list of trusted ports.

Syntax Description

ports	Specifies the trusted ports list.
ports-to-add	Specifies the ports to add to the trusted ports list.
all	Specifies all the ports.
ports-to-delete	Specifies the ports to delete from the trusted ports list.

Default

N/A.

Usage Guidelines

Traffic from trusted ports will be ignored when DoS protect counts the packets to the CPU. If we know that a machine connected to a certain port on the switch is a safe "trusted" machine, and we know that we will not get a DoS attack from that machine, the port where this machine is connected to can be configured as a trusted port, even though a large amount of traffic is going through this port.

Example

This example sets the trusted port list to 3:1-3:7:

```
configure dos-protect trusted-ports ports 3:1-3:7
```

This example adds the trusted port 3:8 to the current list (use this command with a network administrator machine not connected to the internet that is attached to port 3:8):

```
configure dos-protect trusted-ports add-ports 3:8
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect type l3-protect alert-threshold

```
configure dos-protect type l3-protect alert-threshold <packets>
```

Description

Configures the denial of service protection alert threshold.

Syntax Description

packets	Specifies how many packets in an interval will cause an alert.

Default

The default is 4000 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection alert. When an alert occurs, the packets are analyzed, and a temporary ACL is applied to the switch.

Example

This example sets the alert threshold to 8000 packets:

```
configure dos-protect type l3-protect alert-threshold 8000
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect type l3-protect notify-threshold

```
configure dos-protect type l3-protect notify-threshold <packets>
```

Description

Configures the denial of service protection notification threshold.

Syntax Description

packets	Specifies how many packets in an interval will cause a notification.

Default

The default is 3500 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection notification.

Example

This example sets the notification threshold to 7500 packets:

```
configure dos-protect type l3-protect notify-threshold 7500
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on all platforms.

configure ip-security anomaly-protection icmp ipv4-max-size

```
configure ip-security anomaly-protection icmp ipv4-max-size <size> {slot [
<slot> | all ]}
```

Description

Configures the maximum IPv4 ICMP allowed size.

Syntax Description

size	Specifies the size of the IPv4 ICMP in bytes.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv4 ICMP allowed size. The absolute maximum is 1023 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection icmp ipv6-max-size

```
configure ip-security anomaly-protection icmp ipv6-max-size <size> {slot [
<slot> | all ]}
```

Description

Configures the maximum ipv6 ICMP allowed size.

Syntax Description

size	Specifies the size of the IPv6 ICMP in bytes.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv6 ICMP allowed size. The absolute maximum is 16K bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection notify cache

```
configure ip-security anomaly-protection notify cache <size> {slot [<slot>
| all ]}
```

Description

Configures the size of local notification cache.

Syntax Description

size	Specifies the size of the local notification cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1000 events.

Usage Guidelines

This command configures the size of local notification cache. Cached events are stored in local memory. The range is between 1 and 1000 events per second. If the cache is full, newer events replace older events.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection notify rate limit

```
Configure ip-security anomaly-protection notify rate limit <value> {slot  
[<slot> | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

value	Specifies the period of the rate limit.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1 second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate window` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is from 1 to 100 events.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection notify rate window

```
configure ip-security anomaly-protection notify rate window <value> {slot  
[<slot> | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

value	Specifies the period of the rate limit.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 10 events per second.

Usage Guidelines

This is a paired command with [configure ip-security anomaly-protection notify rate limit](#) that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is between 1 and 300 seconds.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection notify trigger off

```
configure ip-security anomaly-protection notify trigger off <value> {slot [<slot> | all ]}
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

value	Specifies the number of events for the trigger.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger on` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



NOTE

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection notify trigger on

configure ip-security anomaly-protection notify trigger on <value> {slot [<slot> | all]}

Description

Configures an anomaly rate-based notification feature.

Syntax Description

value	Specifies the number of events for the trigger.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger off` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



NOTE

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure ip-security anomaly-protection tcp

```
configure ip-security anomaly-protection tcp min-header-size <size> {slot [
<slot> | all ]}
```

Description

Configures the minimum TCP header allowed.

Syntax Description

size	Specifies the size of the header in bytes.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default value is 20 bytes.

Usage Guidelines

This command configures the minimum TCP header allowed. It takes effect for both IPv4 and IPv6 TCP packets.

The range of the minimum TCP header may be between 8 and 255 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

configure mac-lockdown-timeout ports aging-time

```
configure mac-lockdown-timeout ports [all | <port_list>] aging-time
<seconds>
```

Description

Configures the MAC address lock down timeout value in seconds for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.
seconds	Configures the length of the time out value in seconds. The default is 15 seconds; the range is 15 to 2,000,000 seconds.

Default

The default is 15 seconds.

Usage Guidelines

This timer overrides the FDB aging time.

This command only sets the duration of the MAC address lock down timer. To enable the lock down timeout feature, use the following command:

```
enable mac-lockdown-timeout ports [all | <port_list>]
```

Example

The following command configures the MAC address lock down timer duration for 300 seconds for ports 2:3, 2:4, and 2:6:

```
configure mac-lockdown-timeout ports 2:3, 2:4, 2:6 aging-time 300
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure ports limit-learning lock-learning

```
configure ports <portlist> vlan <vlan_name> [limit-learning <number> |
lock-learning | unlimited-learning | unlock-learning]
```

Description

Configures virtual ports for limited or locked MAC address learning.

Syntax Description

portlist	Specifies one or more ports or slots and ports.
vlan_name	Specifies the name of the VLAN.
limit-learning <number>	Specifies a limit on the number of MAC addresses that can be dynamically learned on the specified ports.
lock-learning	Specifies that the current FDB entries for the specified ports should be made permanent static, and no additional learning should be allowed.
unlimited-learning	Specifies that there should not be a limit on MAC addresses that can be learned.
unlock-learning	Specifies that the port should be unlocked (allow unlimited, dynamic learning).

Default

Unlimited, unlocked learning.

Usage Guidelines

If you have enabled ESRP, see the *ExtremeXOS Concepts Guide* for information about using this feature with ESRP.

Limited learning. The limited learning feature allows you to limit the number of dynamically-learned MAC addresses per VLAN. When the learned limit is reached, all new source MAC addresses are blackholed at both the ingress and egress points. This prevent these MAC addresses from learning and responding to Internet control message protocol (ICMP) and address resolution protocol (ARP) packets.

If the limit you configure is greater than the current number of learned entries, all the current learned entries are purged.

Dynamically learned entries still get aged, and can be cleared. If entries are cleared or aged out after the learning limit has been reached, new entries will then be able to be learned until the limit is reached again.

Permanent static and permanent dynamic entries can still be added and deleted using the `create fdbentry` and `delete fdbentry` commands. These override any dynamically learned entries.

For ports that have a learning limit in place, the following traffic still flows to the port:

- Packets destined for permanent MACs and other non-blackholed MACs
- Broadcast traffic
- EDP traffic

Traffic from the permanent MAC and any other non-blackholed MACs will still flow from the virtual port.

If you configure a MAC address limit on VLANs that participate in an Extreme Standby Router Protocol (ESRP) domain, you should add an additional back-to-back link (that has no MAC address limit on these ports) between the ESRP-enabled switches. Doing so prevents ESRP protocol data units (PDUs) from being dropped due to MAC address limit settings.

Port lockdown. The port lockdown feature allows you to prevent any additional learning on the virtual port, keeping existing learned entries intact. This is equivalent to making the dynamically-learned entries permanent static, and setting the learning limit to zero. All new source MAC addresses are blackholed.

Locked entries do not get aged, but can be deleted like any other permanent FDB entries. The maximum number of permanent lockdown entries is 1024. Any FDB entries above will be flushed and blackholed during lockdown.

For ports that have lockdown in effect, the following traffic still flows to the port:

- Packets destined for the permanent MAC and other non-blackholed MACs
- Broadcast traffic
- EDP traffic

Traffic from the permanent MAC will still flow from the virtual port.

Once the port is locked down, all the entries become permanent and will be saved across reboot.

When you remove the lockdown using the unlock-learning option, the learning-limit is reset to unlimited, and all associated entries in the FDB are flushed.

To display the locked entries on the switch, use the following command:

```
show fdb
```

Locked MAC address entries have the “l” flag.

To verify the MAC security configuration for the specified VLAN or ports, use the following commands:

```
show vlan <vlan name> security
show ports <portlist> info detail
```

Example

The following command limits the number of MAC addresses that can be learned on ports 1, 2, 3, and 6 in a VLAN named *accounting*, to 128 addresses:

```
configure ports 1, 2, 3, 6 vlan accounting learning-limit 128
```

The following command locks ports 4 and 5 of VLAN *accounting*, converting any FDB entries to static entries, and prevents any additional address learning on these ports:

```
configure ports 4,5 vlan accounting lock-learning
```

The following command removes the learning limit from the specified ports:

```
configure ports 1, 2, vlan accounting unlimited-learning
```

The following command unlocks the FDB entries for the specified ports:

```
configure ports 4,5 vlan accounting unlock-learning
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure radius server

```
configure radius {mgmt-access | netlogin} [primary | secondary] server
[<ipaddress> | <hostname>] {<udp_port>} client-ip [<ipaddress>] {vr
<vr_name>}
```

Description

Configures the primary and secondary RADIUS authentication server.

Syntax Description

mgmt-access	Specifies the RADIUS authentication server for switch management.
netlogin	Specifies the RADIUS authentication server for network login.
primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
ipaddress	The IP address of the server being configured.
hostname	The host name of the server being configured.
udp_port	The UDP port to use to contact the RADIUS authentication server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.
vr_name	Specifies the virtual router on which the client IP is located. NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The following lists the default behavior of this command:

- The UDP port setting is 1812
- The virtual router used is VR-Mgmt, the management virtual router
- Switch management and network login use the same primary and secondary RADIUS servers for authentication.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Use this command to specify RADIUS server information.

Use of the <hostname> parameter requires that DNS be enabled.

The RADIUS server defined by this command is used for user name authentication and CLI command authentication.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS authentication servers for switch management and another pair for network login. To specify RADIUS authentication servers for switch management (Telnet, SSH, and console sessions), use the `mgmt-access` keyword. To specify RADIUS authentication servers for network login, use the `netlogin` keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS authentication servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS server configuration. Both switch management and network login use the RADIUS authentication server specified in the older configuration.

Example

The following command configures the primary RADIUS server on host `radius1` using the default UDP port (1812) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of `VR-Default`:

```
configure radius primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following command configures the primary RADIUS server for network login authentication on host `netlog1` using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using, by default, the management virtual router interface:

```
configure radius netlogin primary server netlog1 client-ip 10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius shared-secret

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-  
secret {encrypted} <string>
```

Description

Configures the authentication string used to communicate with the RADIUS authentication server.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Configures the authentication string for the primary RADIUS server.
secondary	Configures the authentication string for the secondary RADIUS server.
encrypted	Indicates that the string is already encrypted.
string	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS server.

The RADIUS server must first be configured for use with the switch as a RADIUS client.

The `mgmt-access` keyword specifies the RADIUS server used for switch management authentication.

The `netlogin` keyword specifies the RADIUS server used for network login authentication.

If you do not specify the `mgmt-access` or `netlogin` keywords, the secret applies to both the primary or secondary switch management and netlogin RADIUS servers.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary RADIUS server for both switch management and network login:

```
configure radius primary shared-secret purplegreen
```

The following command configures the shared secret as “redblue” on the primary switch management RADIUS server:

```
configure radius mgmt-access primary shared-secret redblue
```

History

This command was first available in ExtremeXOS 10.1.

The `encrypted` keyword was added in ExtremeXOS 11.0.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius timeout

```
configure radius {mgmt-access | netlogin} timeout <seconds>
```

Description

Configures the timeout interval for RADIUS authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
seconds	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used. After six failed attempts, local user authentication will be used.

The `mgmt-access` keyword specifies the RADIUS server used for switch management authentication.

The `netlogin` keyword specifies the RADIUS server used for network login authentication.

If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS servers.

Example

The following command configures the timeout interval for RADIUS authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used. After 60 seconds (six attempts) local user authentication is used.

```
configure radius timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius-accounting server

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary]
server [<ipaddress> | <hostname>] {<tcp_port>} client-ip [<ipaddress>] {vr
<vr_name>}
```

Description

Configures the RADIUS accounting server.

Syntax Description

mgmt-access	Specifies the RADIUS accounting server for switch management.
netlogin	Specifies the RADIUS accounting server for network login.
primary	Configure the primary RADIUS accounting server.
secondary	Configure the secondary RADIUS accounting server.
ipaddress	The IP address of the accounting server being configured.
hostname	The host name of the accounting server being configured.
tcp_port	The UDP port to use to contact the RADIUS accounting server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS accounting server.
vr_name	Specifies the virtual router on which the client IP is located.
	NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The following lists the default behavior of this command:

- The UDP port setting is 1813
- The virtual router used is VR-Mgmt, the management virtual router
- Switch management and network login use the same RADIUS accounting server.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Use this command to specify the radius accounting server.

The accounting server and the RADIUS authentication server can be the same.

Use of the <hostname> parameter requires that DNS be enabled.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS accounting servers for switch management and another pair for network login. To specify RADIUS accounting servers for switch management (Telnet, SSH, and console sessions), use the `mgmt-access` keyword. To specify RADIUS accounting servers for network login, use the `netlogin` keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS accounting servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS accounting server configuration. Both switch management and network login use the RADIUS accounting server specified in the older configuration.

Example

The following command configures RADIUS accounting on host `radius1` using the default UDP port (1813) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of `VR-Default` for both management and network login:

```
configure radius-accounting primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following command configures RADIUS accounting for network login on host `netlog1` using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using the default virtual router interface:

```
configure radius-accounting netlogin primary server netlog1 client-ip 10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius-accounting shared-secret

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary]
shared-secret {encrypted} <string>
```

Description

Configures the authentication string used to communicate with the RADIUS accounting server.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
primary	Configures the authentication string for the primary RADIUS accounting server.
secondary	Configures the authentication string for the secondary RADIUS accounting server.
encrypted	Indicates that the string is already encrypted.
string	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS accounting server.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the secret applies to both the primary or secondary switch management and `netlogin` RADIUS accounting servers.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purpleaccount” on the primary RADIUS accounting server for both management and network login:

```
configure radius primary shared-secret purpleaccount
```

The following command configures the shared secret as “greenaccount” on the primary management RADIUS accounting server:

```
configure radius mgmt-access primary shared-secret greenaccount
```

History

This command was first available in ExtremeXOS 10.1.

The `encrypted` keyword was added in ExtremeXOS 11.0.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius-accounting timeout

```
configure radius-accounting {mgmt-access | netlogin} timeout <seconds>
```

Description

Configures the timeout interval for RADIUS-Accounting authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
seconds	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds.

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS-Accounting authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS accounting servers.

Example

This example configures the timeout interval for RADIUS-Accounting authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used:

```
configure radius-accounting timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ssh2 key

```
configure ssh2 key {pregenerated}
```

Description

Generates the Secure Shell 2 (SSH2) host key.

Syntax Description

pregenerated	Indicates that the SSH2 authentication key has already been generated. The user will be prompted to enter the existing key.

Default

The switch generates a key for each SSH2 session.

Usage Guidelines

Secure Shell 2 (SSH2) is a feature of ExtremeXOS that allows you to encrypt session data between a network administrator using SSH2 client software and the switch or to send encrypted data from the switch to an SSH2 client on a remote system. Configuration, policy, image, and public key files may also be transferred to the switch using the Secure Copy Program 2 (SCP2)

SSH2 functionality is not present in the base ExtremeXOS software image, but is available as an additional, installable module. Before you can access any SSH2 commands, you must install the module. Without the module, the SSH2 commands do not appear on the command line. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options."](#)

After you have installed the SSH2 module, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

An authentication key must be generated before the switch can accept incoming SSH2 sessions. This can be done automatically by the switch, or you can enter a previously generated key.

If you elect to have the key generated, the key generation process can take up to ten minutes, and cannot be canceled after it has started. Once the key has been generated, you should save your configuration to preserve the key.

To use a key that has been previously created, use the `pregenerated` keyword. You are prompted to enter the pregenerated key. You can use the `show configuration` command to list the previously generated key then copy and paste it after the `configure ssh2 key {pregenerated}` command.



NOTE

Keys generated by ExtremeXOS cannot be used on switches running ExtremeWare images, and keys generated by ExtremeWare cannot be used on switches running ExtremeXOS images.

The key generation process generates the SSH2 private host key. The SSH2 public host key is derived from the private host key, and is automatically transmitted to the SSH2 client at the beginning of an SSH2 session.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions, whether a valid key is present, and the TCP port and virtual router that is being used.

Example

The following command generates an authentication key for the SSH2 session:

```
configure ssh2 key
```

The command responds with the following messages:

```
WARNING: Generating new server host key
This will take approximately 10 minutes and cannot be canceled.
Continue? (y/n)
```

If you respond yes, the command begins the process.

To configure an SSH2 session using a previously generated key, use the following command:

```
configure ssh2 key pregenerated <pre-generated key>
```

Enter the previously-generated key (you can copy and paste it from the saved configuration file; a part of the key pattern is similar to 2d:2d:2d:2d:20:42:45:47:).

History

This command was first available in the ExtremeXOS 11.0 SSH module.

Platform Availability

This command is available on all platforms.

configure sshd2 user-key add user

```
configure sshd2 user-key <key_name> add user <user_name>
```

Description

Associates a user to a key.

Syntax Description

key_name	Specifies the name of the public key.
user_name	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command associates (or *binds*) a user to a key.

Example

The following example binds the key `id_dsa_2048` to user `admin`.

```
configure sshd2 user-key id_dsa_2048 add user admin
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure sshd2 user-key delete user

```
configure sshd2 user-key <key_name> delete user <user_name>
```

Description

Disassociates a user to a key.

Syntax Description

key_name	Specifies the name of the public key.
user_name	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command disassociates (or *unbinds*) a user to a key.

Example

The following example unbinds the key `id_dsa_2048` from user `admin`.

```
configure sshd2 user-key id_dsa_2048 delete user admin
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure ssl certificate pregenerated

```
configure ssl certificate pregenerated
```

Description

Obtains the pre-generated certificate from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must upload or generate a certificate for SSL server use. With this command, you copy and paste the certificate into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: MD5 and SHA.

This command is also used when downloading or uploading the configuration. Do not modify the certificate stored in the uploaded configuration file because the certificate is signed using the issuer's private key.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

The Converged Network Analyzer (CNA) Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [Appendix C, "CNA Agent."](#)

Example

The following command obtains the pre-generated certificate from the user:

```
configure ssl certificate pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a blank line to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure ssl certificate privkeylen

```
configure ssl certificate privkeylen <length> country <code> organization
<org_name> common-name <name>
```

Description

Creates a self signed certificate and private key that can be saved in the EEPROM.

Syntax Description

length	Specifies the private key length in bytes. Valid values are between 1024 and 4096.
code	Specifies the country code in 2-character form.
org_name	Specifies the organization name. The organization name can be up to 64 characters long.
name	Specifies the common name. The common name can be up to 64 characters long.

Default

N/A.

Usage Guidelines

This command creates a self signed certificate and private key that can be saved in the EEPROM. The certificate generated is in the PEM format.

Any existing certificate and private key is overwritten.

The size of the certificate depends on the RSA key length (`privkeylen`) and the length of the other parameters (`country`, `organization` name, and so forth) supplied by the user. If the RSA key length is 1024, then the certificate is approximately 1 kb. For an RSA key length of 4096, the certificate length is approximately 2 kb, and the private key length is approximately 3 kb.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (`ssh.xmod`). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

The CNA Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [Appendix C, "CNA Agent Commands."](#)

Example

The following command creates an SSL certificate in the USA for a website called bigcats:

```
configure ssl certificate privkeylen 2048 country US organization IEEE common-name  
bigcats
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure ssl privkey pregenerated

```
configure ssl privkey pregenerated
```

Description

Obtains the pre-generated private key from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command is also used when downloading or uploading the configuration. The private key is stored in the EEPROM, and the certificate is stored in the configuration file.

With this command, you copy and paste the private key into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: MD5 and SHA.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

The CNA Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [Appendix C, "CNA Agent Commands."](#)

Example

The following command obtains the pre-generated private key from the user:

```
configure ssl privkey pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a blank line to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure tacacs server

```
configure tacacs [primary | secondary] server [<ipaddress> | <hostname>]
{<tcp_port>} client-ip <ipaddress> {vr <vr_name>}
```

Description

Configures the server information for a TACACS+ authentication server.

Syntax Description

primary	Configures the primary TACACS+ server.
secondary	Configures the secondary TACACS+ server.
ipaddress	The IP address of the TACACS+ server being configured.
hostname	The host name of the TACACS+ server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ server.
vr_name	Specifies the virtual router on which the client IP is located. NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

TACACS+ uses TCP port 49. The default virtual router is VR-Mgmt, the management virtual router

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Use this command to configure the server information for a TACACS+ server.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command configures server tacacs1 as the primary TACACS+ server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure tacacs shared-secret

```
configure tacacs [primary | secondary] shared-secret {encrypted} <string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ authentication server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ server.
secondary	Configures the authentication string for the secondary TACACS+ server.
encrypted	Indicates that the string is already encrypted.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

The secret must be the same between the client switch and the TACACS+ server.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary TACACS+ server:

```
configure tacacs-accounting primary shared-secret purplegreen
```

History

This command was first available in ExtremeXOS 10.1.

The `encrypted` keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure tacacs timeout

```
configure tacacs timeout <seconds>
```

Description

Configures the timeout interval for TACAS+ authentication requests.

Syntax Description

seconds	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds.

Default

The default is 3 seconds.

Usage Guidelines

Use this command to configure the timeout interval for TACACS+ authentication requests.

To detect and recover from a TACACS+ server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it will take 3 seconds to fail over from the primary TACACS+ server to the secondary TACACS+ server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ authentication to 10 seconds:

```
configure tacacs timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure tacacs-accounting server

```
configure tacacs-accounting [primary | secondary] server [<ipaddress> |
<hostname>] {<udp_port>} client-ip <ipaddress> {vr <vr_name>}
```

Description

Configures the TACACS+ accounting server.

Syntax Description

primary	Configures the primary TACACS+ accounting server.
secondary	Configures the secondary TACACS+ accounting server.
ipaddress	The IP address of the TACACS+ accounting server being configured.
hostname	The host name of the TACACS+ accounting server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ accounting server.
vr_name	Specifies the virtual router on which the client IP is located. NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

Unconfigured. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

You can use the same TACACS+ server for accounting and authentication.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Example

The following command configures server tacacs1 as the primary TACACS+ accounting server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs-accounting primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure tacacs-accounting shared-secret

```
configure tacacs-accounting [primary | secondary] shared-secret {encrypted}  
<string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ accounting server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ accounting server.
secondary	Configures the authentication string for the secondary TACACS+ accounting server.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

Secret needs to be the same as on the TACACS+ server.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “tacacsaccount” on the primary TACACS+ accounting server:

```
configure tacacs-accounting primary shared-secret tacacsaccount
```

History

This command was first available in ExtremeXOS 10.1.

The `encrypted` keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure tacacs-accounting timeout

```
configure tacacs-accounting timeout <seconds>
```

Description

Configures the timeout interval for TACACS+ accounting authentication requests.

Syntax Description

seconds	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for TACACS+ accounting authentication requests.

To detect and recover from a TACACS+ accounting server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ accounting server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ accounting server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the accounting server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it takes 3 seconds to fail over from the primary TACACS+ accounting server to the secondary TACACS+ accounting server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ accounting authentication to 10 seconds:

```
configure tacacs-accounting timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure trusted-ports trust-for dhcp-server

```
configure trusted-ports [<ports>|all] trust-for dhcp-server
```

Description

Configures one or more trusted DHCP ports.

Syntax Description

ports	Specifies one or more ports to be configured as trusted ports.
all	Specifies all ports to be configured as trusted ports.

Default

N/A.

Usage Guidelines

To configure trusted DHCP ports, you must first enable DHCP snooping on the switch. To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>] violation-  
action [drop-packet {[block-mac | block-port] [duration <duration_in_seconds> |  
permanently] | none}}] {snmp-trap}
```

Trusted ports do not block traffic; rather, the switch forwards any DHCP server packets that appear on trusted ports. Depending on your DHCP snooping configuration, the switch drops packets and can disable the port temporarily, disable the port permanently, blackhole the MAC address temporarily, blackhole the MAC address permanently, and so on.

If you configure one or more trusted ports, the switch assumes that all DHCP server packets on the trusted port are valid.

Displaying DHCP Trusted Server Information. To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command configures ports 2:2 and 2:3 as trusted ports:

```
configure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure trusted-servers vlan add server trust-for dhcp-server

```
configure trusted-servers {vlan} <vlan_name> add server <ip_address> trust-  
for dhcp-server
```

Description

Configures and enables a trusted DHCP server on the switch.

Syntax Description

vlan_name	Specifies the VLAN name.
ip_address	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

If you configured trusted DHCP server, the switch forwards only DHCP packets from the trusted servers. The switch drops DHCP packets from other DHCP snooping-enabled ports.

You can configure a maximum of eight trusted DHCP servers on the switch.

If you configure a port as a trusted port, the switch assumes that all DHCP server packets on that port are valid.

Displaying DHCP Trusted Server Information. To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command configures a trusted DHCP server on the switch:

```
configure trusted-servers vlan purple add server 10.10.10.10 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure trusted-servers vlan delete server trust-for dhcp-server

```
configure trusted-servers vlan <vlan_name> delete server <ip_address>
trust-for dhcp-server
```

Description

Deletes a trusted DHCP server from the switch.

Syntax Description

vlan_name	Specifies the VLAN name.
ip_address	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

Use this command to delete a trusted DHCP server from the switch.

Displaying DHCP Trusted Server Information. To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command deletes a trusted DHCP server from the switch:

```
configure trusted-servers vlan purple delete server 10.10.10.10 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure vlan dhcp-address-range

```
configure vlan <vlan_name> dhcp-address-range <ipaddress1> - <ipaddress2>
```

Description

Configures a set of DHCP addresses for a VLAN.

Syntax Description

vlan_name	Specifies the VLAN on whose ports DHCP will be enabled.
ipaddress1	Specifies the first IP address in the DHCP address range to be assigned to this VLAN.
ipaddress2	Specifies the last IP address in the DHCP address range to be assigned to this VLAN.

Default

N/A.

Usage Guidelines

The following error conditions are checked: ipaddress2 >= ipaddress1, the range must be in the VLAN's network, the range does not contain the VLAN's IP address, and the VLAN has an IP address assigned.

Example

The following command allocates the IP addresses between 192.168.0.20 and 192.168.0.100 for use by the VLAN *temporary*:

```
configure temporary dhcp-address-range 192.168.0.20 - 192.168.0.100
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan dhcp-lease-timer

```
configure vlan <vlan_name> dhcp-lease-timer <lease-timer>
```

Description

Configures the timer value in seconds returned as part of the DHCP response.

Syntax Description

name	Specifies the VLAN on whose ports netlogin should be disabled.
lease-timer	Specifies the timer value, in seconds.

Default

N/A.

Usage Guidelines

The timer value is specified in seconds. The timer value range is 0 - 4294967295, where 0 indicates the default (not configured) value of 7200 second.

Example

The following command configures the DHCP lease timer value for VLAN *corp*:

```
configure vlan corp dhcp-lease-timer <lease-timer>
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan dhcp-options

```
configure vlan <vlan_name> dhcp-options [default-gateway | dns-server |
wins-server] <ipaddress>
```

Description

Configures the DHCP options returned as part of the DHCP response by a switch configured as a DHCP server.

Syntax Description

name	Specifies a VLAN name.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
wins-server	Specifies the NetBIOS name server (NBNS) option.
ipaddress	The IP address associated with the specified option.

Default

N/A.

Usage Guidelines

This command configures the DHCP options that can be returned to the DHCP client. For the `default-gateway` option you are only allowed to configure an IP address that is in the VLAN's network range. For the other options, any IP address is allowed.

The options below represent the following BOOTP options specified by RFC2132:

- `default-gateway`—Router option, number 3
- `dns-server`—Domain Name Server option, number 6
- `wins-server`—NetBIOS over TCP/IP Name Server option, number 44

Example

The following command configures the DHCP server to return the IP address 10.10.20.8 as the router option:

```
configure vlan <name> dhcp-options default-gateway 10.10.20.8
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create sshd2 key-file

```
create sshd2 key-file {host-key | user-key} <key_name>
```

Description

Creates a file for the user-key or host-key.

Syntax Description

host-key	Specifies the name of the host-key
user-key	Specifies the name of the user-key.
key_name	Specifies the name of the public key.

Default

N/A.

Usage Guidelines

This command is used to write the user or the host public key in a file. The key files will be created with a .ssh file extension; this enables the administrator to copy the public key files to another server.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

create sshd2 user-key

```
create sshd2 user-key <key_name> <key> {subject <subject>} {comment  
<comment>}
```

Description

Creates a user key.

Syntax Description

key_name	Specifies the name of the public key.
key	Specifies the key. Note: The key cannot have any spaces in it.
subject	Specifies the subject.
comment	Specifies the comment (an optional field)

Default

N/A.

Usage Guidelines

This command is used to enter, or cut and paste, your public key. You can also enter the public key into the switch by using the SCP or SFTP client that is connected to the switch.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

delete sshd2 user-key

```
delete sshd2 user-key <key_name>
```

Description

Deletes a user key.

Syntax Description

key_name	Specifies the name of the public key to be deleted.

Default

N/A.

Usage Guidelines

This command is used to delete a user key. The key is deleted regardless of whether or not it is bound to a user.



NOTE

If a user is bound to the key, they are first unbound or unassociated, and then the key is deleted

Example

The following example shows the SSH user key `id_dsa_2048` being deleted:

```
delete sshd2 user-key id_dsa_2048
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable dhcp ports vlan

```
disable dhcp ports <portlist> vlan <vlan_name>
```

Description

Disables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be disabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be disabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables DHCP for port 6:9 in VLAN *corp*:

```
disable dhcp ports 6:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable dos-protect

```
disable dos-protect
```

Description

Disables denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command disables denial of service protection.

```
disable dos-protect
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable iparp gratuitous protect

```
disable iparp gratuitous protect vlan <vlan-name>
```

Description

Disables gratuitous ARP protection on the specified VLAN.

Syntax Description

vlan-name	Specifies the VLAN.

Default

Disabled.

Usage Guidelines

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network. Also, the switch will add a MAC ACL to blackhole the offending host.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN *corp*:

```
disable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ip-security anomaly-protection

```
disable ip-security anomaly-protection {slot [ <slot> | all ]}
```

Description

Disables all anomaly checking options.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This commands disables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection ip

```
disable ip-security anomaly-protection ip { slot [ <slot> | all ] }
```

Description

Disables source and destination IP address checking.

Syntax Description

slot	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection l4port

```
disable ip-security anomaly-protection l4port [tcp | udp | both] {slot [
<slot> | all ]}
```

Description

Disables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be disabled for checking.
udp	Specifies that the UDP port be disabled for checking.
both	Specifies both the TCP and UDP ports be disabled for checking.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a layer 4 protocol error. (This type of error can be found in a BALT attack.)



NOTE

Options udp and tcp are supported on Summit X250 platform

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection tcp flags

```
disable ip-security anomaly-protection tcp flags {slot [ <slot> | all ]}
```

Description

Disables TCP flag checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection tcp fragment

```
disable ip-security anomaly-protection tcp fragment {slot [ <slot> | all ]}
```

Description

Disables TCP fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size
- If its IP offset field==1 (for IPv4 only)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection icmp

```
disable ip-security anomaly-protection icmp {slot [ <slot> | all ]}
```

Description

Disables ICMP size and fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets for IPv4 packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security anomaly-protection notify

```
disable ip-security anomaly-protection notify [log | snmp | cache] {slot [
<slot> | all ]}
```

Description

Disables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- **log:** log anomaly events in the switch log system; you can view and manage this log with the `show log` and `configure log` commands
- **snmp:** the anomaly events generate SNMP traps
- **cache:** logs the most recent and unique anomaly events in memory; rebooting the switch will cause all the logged events to be lost (the number of cached events is configured by command)

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

disable ip-security arp gratuitous-protection vlan

```
disable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]
```

Description

Disables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
vlan-name	Specifies the VLAN.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the [disable iparp gratuitous protect](#) command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network. Also, the switch will add a MAC ACL to blackhole the offending host.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN *corp*:

```
disable ip-security arp gratuitous-protection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable ip-security arp learning learn-from-arp vlan ports

```
disable ip-security arp learning learn-from-arp {vlan} <vlan_name> ports
[all | <ports>]
```

Description

Disables ARP learning on the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

You can disable ARP learning so that the only entries in the ARP table are either manually added or those created by DHCP secured ARP; the switch does not add entries by tracking ARP requests and replies. By disabling ARP learning and adding a permanent entry or configuring DHCP secured ARP, you can centrally manage and allocate client IP addresses and prevent duplicate IP addresses from interrupting network operation.

To manually add a permanent entry to the ARP table, use the following command:

```
configure iparp add <ip_addr> {vr <vr_name>} <mac>
```

To configure DHCP secure ARP as a method to add entries to the ARP table, use the following command:

```
enable ip-security arp learning learn-from-dhcp vlan <vlan_name> ports [all | <ports>]
{poll-interval <interval_in_seconds>} {retries <number_of_retries>}
```

Displaying ARP Information. To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

**NOTE**

DHCP secured ARP entries are stored as static entries in the ARP table.

Example

The following command disables ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable ip-security arp learning learn-from-dhcp vlan ports

```
disable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports  
[all | <ports>]
```

Description

Disables DHCP secured ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to disable DHCP secured ARP learning.

Displaying ARP Information. To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command disables DHCP secured ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable ip-security arp validation vlan

```
disable ip-security arp validation {vlan} <vlan_name> [all | <ports>]
```

Description

Disables ARP validation for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports.
ports	Specifies one or more ports.

Default

By default, ARP validation is disabled.

Usage Guidelines

Use this command to disable ARP validation.

Displaying ARP Validation Information. To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} <vlan_name>
```

Example

The following command disables ARP validation on port 1:1 of the VLAN valid:

```
disable ip-security arp validation vlan valid ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable ip-security dhcp-snooping

```
disable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]
```

Description

Disables DHCP snooping on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.
all	Specifies all ports to stop receiving DHCP packets.
ports	Specifies one or more ports to stop receiving DHCP packets.

Default

By default, DHCP snooping is disabled

Usage Guidelines

Use this command to disable DHCP snooping on the switch.

Example

The following command disables DHCP snooping on the switch:

```
disable ip-security dhcp-snooping vlan snoop ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable ip-security source-ip-lockdown ports

```
disable ip-security source-ip-lockdown ports [all | <ports>]
```

Description

Disables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be disabled.
ports	Specifies one or more ports for which source IP lockdown should be disabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command disables source IP lockdown on ports 1:1 and 1:4:

```
disable ip-security source-ip-lockdown ports 1:1, 1:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable mac-lockdown-timeout ports

```
disable mac-lockdown-timeout ports [all | <port_list>]
```

Description

Disables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down feature is disabled.

Usage Guidelines

If you disable the MAC lock down timer on a port, existing MAC address entries for the port will time out based on the FDB aging period.

Example

The following command disables the MAC address lock down timer set for ports 2:3 and 2:4:

```
disable mac-lockdown-timeout ports 2:3, 2:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable radius

```
disable radius {mgmt-access | netlogin}
```

Description

Disables the RADIUS client.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Use the `mgmt-access` keyword to disable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to disable RADIUS authentication for network login.

If you do not specify a keyword, RADIUS authentication is disabled on the switch for both management and network login.

Example

The following command disables RADIUS authentication on the switch for both management and network login:

```
disable radius
```

The following command disables RADIUS authentication on the switch for network login:

```
disable radius netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable radius-accounting

```
disable radius-accounting {mgmt-access | netlogin}
```

Description

Disables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

Use the `mgmt-access` keyword to disable RADIUS accounting for switch management functions.

Use the `netlogin` keyword to disable RADIUS accounting for network login.

If you do not specify a keyword, RADIUS accounting is disabled on the switch for both management and network login.

Example

The following command disables RADIUS accounting on the switch for both management and network login:

```
disable radius-accounting
```

The following command disables RADIUS accounting on the switch for network login:

```
disable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ssh2

```
disable ssh2
```

Description

Disables the SSH2 server for incoming SSH2 sessions to switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

SSH2 options (non-default port setting) are not saved when SSH2 is disabled.

To view the status of SSH2 on the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enable/disable state for SSH2.

Example

The following command disables the SSH2 server:

```
disable ssh2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs

```
disable tacacs
```

Description

Disables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ authentication for the switch:

```
disable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs-accounting

```
disable tacacs-accounting
```

Description

Disables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ accounting:

```
disable tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs-authorization

```
disable tacacs-authorization
```

Description

Disables TACACS+ authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disables CLI command authorization but leaves user authentication enabled.

Example

The following command disables TACACS+ CLI command authorization:

```
disable tacacs-authorization
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable web http

```
disable web http
```

Description

Disables the hypertext transfer protocol (HTTP) access to the switch on the default port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to disallow users from connecting with HTTP. Disabling HTTP access forces user to use a secured HTTPS connection if web HTTPS is enabled.

Use the following command to enable web HTTPS:

```
enable web https
```

Example

The following command disables HTTP on the default port:

```
disable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

disable web https

```
disable web https
```

Description

Disables the secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable SSL before changing the certificate or private key.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

Example

The following command disables SSL on the default port:

```
disable web https
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

download ssl certificate

```
download ssl <ip_address> certificate <cert file>
```

Description

Permits downloading of a certificate key from files stored in a TFTP server.

Syntax Description

ip_address	Specifies the IP address of the TFTP server.
cert file	Specifies the name of the certificate key.

Default

N/A.

Usage Guidelines

If the download operation is successful, any existing certificate is overwritten. After a successful download, the software attempts to match the public key in the certificate against the private key stored. If the private and public keys do not match, the switch displays a warning message similar to the following: **Warning: The Private Key does not match with the Public Key in the certificate.** This warning acts as a reminder to also download the private key.



NOTE

You can only download a certificate key in the VR-Mgmt virtual router.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the save command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

You can purchase and obtain SSL certificates from Internet security vendors.

Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Colon (:)

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a certificate from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 certificate g0ethner1
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

download ssl privkey

```
download ssl <ip_address> privkey <key file>
```

Description

Permits downloading of a private key from files stored in a TFTP server.

Syntax Description

ip_address	Specifies the IP address of the TFTP server.
key file	Specifies the name of the private key file.

Default

N/A.

Usage Guidelines

If the operation is successful, the existing private key is overwritten.

After a successful download, a check is performed to find out whether the private key downloaded matches the public key stored in the certificate. If the private and public keys do not match, the switch displays a warning similar to the following: `Warning: The Private Key does not match with the Public Key in the certificate.` This warning acts as a reminder to also download the corresponding certificate.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the `save` command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, "Software Upgrade and Boot Options,"](#) of the *ExtremeXOS Concepts Guide*.

Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)

- Period (.)
- Colon (:)

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a private key from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 privkey t00Ts1e
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

enable dhcp ports vlan

```
enable dhcp ports <portlist> vlan <vlan_name>
```

Description

Enables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be enabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be enabled.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables DHCP for port 5:9 in VLAN *corp*:

```
enable dhcp ports 5:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable dos-protect

```
enable dos-protect
```

Description

Enables denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

None.

Example

The following command enables denial of service protection.

```
enable dos-protect
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable dos-protect simulated

```
enable dos-protect simulated
```

Description

Enables simulated denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

If simulated denial of service is enabled, no ACLs are created. This mode is useful to gather information about normal traffic levels on the switch. This will assist in configuring denial of service protection so that legitimate traffic is not blocked.

Example

The following command enables simulated denial of service protection.

```
enable dos-protect simulated
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

enable iparp gratuitous protect

```
enable iparp gratuitous protect vlan <vlan-name>
```

Description

Enables gratuitous ARP protection on the specified VLAN.

Syntax Description

vlan-name	Specifies the VLAN.

Default

By default, gratuitous ARP is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, the [enable ip-security arp gratuitous-protection vlan](#) command replaces this command for configuring gratuitous ARP.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network. Also, the switch will add a MAC ACL to blackhole the offending host.

Example

The following command enables gratuitous ARP protection for VLAN *corp*:

```
enable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ip-security anomaly-protection

```
enable ip-security anomaly-protection {slot [ <slot> | all ]}
```

Description

Enables all anomaly checking options.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This commands enables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection icmp

```
enable ip-security anomaly-protection icmp {slot [ <slot> | all ]}
```

Description

Enables ICMP size and fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection ip

```
enable ip-security anomaly-protection ip { slot [ <slot> | all ] }
```

Description

Enables source and destination IP address checking.

Syntax Description

slot	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection l4port

```
enable ip-security anomaly-protection l4port [tcp | udp | both] {slot [
<slot> | all ]}
```

Description

Enables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be enabled for checking.
udp	Specifies that the UDP port be enabled for checking.
both	Specifies both the TCP and UDP ports be enabled for checking.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enabled TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a layer 4 protocol error. (This type of error can be found in a BALT attack.)



NOTE

Options udp and tcp are supported on Summit X250 platform

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection notify

```
enable ip-security anomaly-protection notify [log | snmp | cache] {slot [
<slot> | all ]}
```

Description

Enables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- **log:** The anomaly events are logged into EMS log.
- **snmp:** The anomaly events generate SNMP traps.
- **cache:** The most recent and unique anomaly events are stored in memory for review and investigation.

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection tcp flags

```
enable ip-security anomaly-protection tcp flags {slot [ <slot> | all ]}
```

Description

Enables TCP flag checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command Enables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security anomaly-protection tcp fragment

```
enable ip-security anomaly-protection tcp fragment {slot [ <slot> | all ]}
```

Description

Enables TCP fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size
- For the first IPv6 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv6 TCP header allowed size
- If its IP offset field==1 (for IPv4 only)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

enable ip-security arp gratuitous-protection vlan

```
enable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]
```

Description

Enables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
vlan-name	Specifies the VLAN.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the `enable iparp gratuitous protect` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network. Also, the switch will add a MAC ACL to blackhole the offending host.

Beginning with ExtremeXOS 11.6, if you enable both DHCP secured ARP and gratuitous ARP protection, the switch protects its own IP address and those of the hosts that appear as secure entries in the ARP table.

To protect the IP addresses of the hosts that appear as secure entries in the ARP table, use the following commands to enable DHCP snooping, DHCP secured ARP, and gratuitous ARP on the switch:

- `enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>] violation-action [drop-packet [[block-mac | block-port] [duration <duration_in_seconds> | permanently] | none]] {snmp-trap}`
- `enable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports [all | <ports>]`
- `enable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]`

Displaying Gratuitous ARP Information. To display information about gratuitous ARP, use the following command:

```
show ip-security arp gratuitous-protection
```

Example

The following command enables gratuitous ARP protection for VLAN *corp*:

```
enable ip-security arp gratuitous-protection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable ip-security arp learning learn-from-arp vlan ports

```
enable ip-security arp learning learn-from-arp {vlan} <vlan_name> ports  
[all | <ports>]
```

Description

Enables ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

ARP is part of the TCP/IP suite used to associate a device's physical address (MAC address) with its logical address (IP address). The switch broadcasts an ARP request that contains the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted across the network. The switch maintains an ARP table (also known as an ARP cache) that displays each MAC address and its corresponding IP address.

By default, the switch builds its ARP table by tracking ARP requests and replies, which is known as ARP learning.

Displaying ARP Information. To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command enables ARP learning on port 1:1 of the VLAN learn:

```
enable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable ip-security arp learning learn-from-dhcp vlan ports

```
enable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports
[all | <ports>]
```

Description

Enables DHCP secured ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to configure the switch to add the MAC address and its corresponding IP address to the ARP table as a secure ARP entry. The switch does not update secure ARP entries, regardless of the ARP requests and replies seen by the switch. DHCP secured ARP is linked to the “DHCP snooping” feature. The same DHCP bindings database created when you enabled DHCP snooping is also used by DHCP secured ARP to create secure ARP entries. The switch only removes secure ARP entries when the corresponding DHCP entry is removed from the trusted DHCP bindings database.



NOTE

If you enable DHCP secured ARP on the switch, ARP learning continues, which allows insecure entries to be added to the ARP table.

The default ARP timeout (configure `iparp timeout`) and ARP refresh (enable `iparp refresh`) settings do not apply to DHCP secured ARP entries. The switch removes DHCP secured ARP entries upon any DHCP release packet received from the DHCP client.

Displaying ARP Information. To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command enables DHCP secured ARP learning on port 1:1 of the VLAN learn and uses the default polling and retry intervals:

```
enable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable ip-security arp validation vlan violation-action

```
enable ip-security arp validation {destination-mac} {source-mac} {ip}
{vlan} <vlan_name> [all | <ports>] violation-action [drop-packet {[block-
port] [duration <duration_in_seconds> | permanently]]] {snmp-trap}
```

Description

Enables ARP validation for the specified VLAN and member ports.

Syntax Description

destination-mac	Specifies that the switch checks the ARP payload for the MAC destination address in the Ethernet header and the receiver's host address in the ARP response.
source-mac	Specifies that the switch checks ARP requests and responses for the MAC source address in the Ethernet header and the sender's host address in the ARP payload.
ip	Specifies the switch checks the IP address in the ARP payload and compares it to the DHCP bindings database. If the IP address does exist in the DHCP bindings table, the switch verifies that the MAC address is the same as the sender hardware address in the ARP request. If not, the packet is dropped.
vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports to participate in ARP validation.
ports	Specifies one or more ports to participate in ARP validation.
drop-packet	Specifies that the switch drops the invalid ARP packet.
block-port	Indicates that the switch blocks invalid ARP requests on the specified port.
permanently	Specifies the switch to permanently disable the port upon receiving an invalid ARP request.
duration_in_seconds	Specifies the switch to temporarily disable the specified port upon receiving an invalid ARP request. The range is seconds.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, ARP validation is disabled.

Usage Guidelines

The violation action setting determines what action(s) the switch takes when an invalid ARP is received.

Depending on your configuration, the switch uses the following methods to check the validity of incoming ARP packets:

- Drop packet—The switch confirms that the MAC address and its corresponding IP address are in the DHCP binding database built by DHCP snooping. This is the default behavior when you enable ARP validation. If the MAC address and its corresponding IP address are in the DHCP bindings

database, the entry is valid. If the MAC address and its corresponding IP address are not in the DHCP bindings database, the entry is invalid, and the switch drops the ARP packet.

- **IP address**—The switch checks the IP address in the ARP payload. If the switch receives an IP address in the ARP payload that is in the DHCP binding database, the entry is valid. If the switch receives an IP address that is not in the DHCP binding database, for example 255.255.255.255 or an IP multicast address, the entry is invalid or unexpected.
- **Source MAC address**—The switch checks ARP requests and responses for the source MAC address in the Ethernet header and the sender's host address in the ARP payload. If the source MAC address and sender's host address are the same, the entry is valid. If the source MAC source and the sender's host address are different, the entry is invalid.
- **Destination MAC address**—The switch checks the ARP payload for the destination MAC address in the Ethernet header and the receiver's host address. If the destination MAC address and the target's host address are the same, the entry is valid. If the destination MAC address and the target's host address are different, the entry is invalid.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS log filters. For more information about EMS, see the EMS commands in [Chapter 5, "Commands for Status Monitoring and Statistics."](#)

Displaying ARP Validation Information. To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} <vlan_name>
```

Example

The following command enables ARP validation on port 1:1 of the VLAN valid:

```
enable ip-security arp validation vlan valid ports 1:1 drop-packet
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable ip-security dhcp-snooping vlan ports violation-action

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]
violation-action [drop-packet {[block-mac | block-port] [duration
<duration_in_seconds> | permanently] | none}}] {snmp-trap}
```

Description

Enables DHCP snooping for the specified VLAN and ports.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN. Create and configure the VLAN before enabling DHCP snooping.
all	Specifies all ports to receive DHCP packets.
ports	Specifies one or more ports to receive DHCP packets.
drop-packet	Indicates that the switch drop the rogue DHCP packet received on the specified port.
block-mac	Indicates that the switch blocks rogue DHCP packets from the specified MAC address on the specified port. The MAC address is added to the DHCP bindings database.
block-port	Indicates that the switch blocks rogue DHCP packets on the specified port. The port is added to the DHCP bindings database.
duration_in_seconds	Specifies that the switch temporarily disable the specified port upon receiving a rogue DHCP packet. The range is seconds.
permanently	Specifies that the switch to permanently disable the specified port upon receiving a rogue DHCP packet.
none	Specifies that the switch takes no action when receiving a rogue DHCP packet; the switch does not drop the packet.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to enable DHCP snooping on the switch.



NOTE

Snooping IP fragmented DHCP packets is not supported.

The violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port or the IP address of the originating server is not among those of the

configured trusted DHCP servers. The DHCP server packets are DHCP OFFER, ACK and NAK. The following list describes the violation actions:

- **block-mac**—The switch automatically generates an ACL to block the MAC address on that port. The switch does not blackhole that MAC address in the FDB. The switch can either temporarily or permanently block the MAC address.
- **block-port**—The switch blocks all incoming rogue DHCP packets on that port. The switch disables the port either temporarily or permanently to block the traffic on that port.
- **none**—The switch takes no action to drop the rogue DHCP packet or block the port, and so on. In this case, DHCP snooping continues to build and manage the DHCP bindings database and DHCP forwarding will continue in hardware as before.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS log filters. For more information about EMS, see the EMS commands in [Chapter 5, “Commands for Status Monitoring and Statistics.”](#)

Displaying DHCP Snooping Information. To display the DHCP snooping configuration settings, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display the DHCP bindings database, use the following command:

```
show ip-security dhcp-snooping entries {vlan} <vlan_name>
```

To display any violations that occur, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command enables DHCP snooping on the switch and has the switch block DHCP packets from port 1:1:

```
enable ip-security dhcp-snooping vlan snoop ports 1:1 violation-action drop-packet  
block-port
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable ip-security source-ip-lockdown ports

```
enable ip-security source-ip-lockdown ports [all | <ports>]
```

Description

Enables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be enabled.
ports	Specifies one or more ports for which source IP lockdown should be enabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

Source IP lockdown prevents IP address spoofing by automatically placing source IP address filters on specified ports. If configured, source IP lockdown allows only traffic from a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address to enter the network.

To configure source IP lockdown, you must enable DHCP snooping on the ports connected to the DHCP server and DHCP client before you enable source IP lockdown. You must enable source IP lockdown on the ports connected to the DHCP client, not on the ports connected to the DHCP server. The same DHCP bindings database created when you enable DHCP snooping is also used by the source IP lockdown feature to create ACLs that permit traffic from DHCP clients. All other traffic is dropped. In addition, the DHCP snooping violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port.

To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>] violation-  
action [drop-packet {[block-mac | block-port] [duration <duration_in_seconds> |  
permanently] | none}}] {snmp-trap}
```

Displaying Source IP Lockdown Information. To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command enables source IP lockdown on ports 1:1 and 1:4:

```
enable ip-security source-ip-lockdown ports 1:1, 1:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable mac-lockdown-timeout ports

```
enable mac-lockdown-timeout ports [all | <port_list>]
```

Description

Enables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down timeout feature is disabled.

Usage Guidelines

You cannot enable the MAC lock down timer on a port that also has the lock learning feature enabled.

Example

The following command enables the MAC address lock down timeout feature for ports 2:3, 2:4, and 2:6:

```
enable mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable radius

```
enable radius {mgmt-access | netlogin}
```

Description

Enables the RADIUS client on the switch.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Before you enable RADIUS on the switch, you must configure the servers used for authentication and configure the authentication string (shared secret) used to communicate with the RADIUS authentication server.

To configure the RADIUS authentication servers, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] server [<ipaddress> | <hostname>] [<udp_port>] client-ip [<ipaddress>] {vr <vr_name>}
```

To configure the shared secret, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-secret {encrypted} <string>
```

If you do not specify a keyword, RADIUS authentication is enabled on the switch for both management and network login. When enabled, all web, Telnet, and SSH logins are sent to the RADIUS servers for authentication. When used with a RADIUS server that supports ExtremeXOS CLI authorization, each CLI command is sent to the RADIUS server for authorization before it is executed.

Use the `mgmt-access` keyword to enable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to enable RADIUS authentication for network login.

Example

The following command enables RADIUS authentication on the switch for both management and network login:

```
enable radius
```

The following command enables RADIUS authentication on the switch for network login:

```
enable radius netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable radius-accounting

```
enable radius-accounting {mgmt-access | netlogin}
```

Description

Enables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

The RADIUS client must also be enabled.

Before you enable RADIUS accounting on the switch, you must configure the servers used for accounting and configure the authentication string (shared secret) used to communicate with the RADIUS accounting server.

To configure the RADIUS accounting servers, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] server
[<ipaddress> | <hostname>] [<tcp_port>] client-ip [<ipaddress>] {vr <vr_name>}
```

To configure the shared secret, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] shared-
secret {encrypted} <string>
```

If you do not specify a keyword, RADIUS accounting is enabled on the switch for both management and network login.

Use the `mgmt-access` keyword to enable RADIUS accounting for switch management functions.

Use the `netlogin` keyword to enable RADIUS accounting for network login.

Example

The following command enables RADIUS accounting on the switch for both management and network login:

```
enable radius-accounting
```

The following command enables RADIUS accounting for network login:

```
enable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ssh2

```
enable ssh2 {access-profile [<access_profile> | none]} {port
<tcp_port_number>} {vr [<vr_name> | all | default]}
```

Description

Enables SSH2 server to accept incoming sessions from SSH2 clients.

Syntax Description

access_profile	Specifies an ACL policy.
none	Cancels a previously configured ACL policy.
port	Specifies a TCP port number. The default is port 22.
vr_name	Specifies a virtual router name.
	NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.
all	Specifies that SSH is enabled on all virtual routers.
default	Specifies that SSH is enabled on the default virtual router.

Default

The SSH2 feature is disabled by default.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

SSH2 enables the encryption of session data. You must be logged in as an administrator to enable SSH2.

SSH2 functionality is not present in the base ExtremeXOS software image, but is in an additional, installable module. Before you can access any SSH2 commands, you must install the module. Without the module, the commands do not appear on the command line. To install the module, see the instructions in [Appendix A, “Software Upgrade and Boot Options.”](#)

After you have installed the SSH2 module, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

Use the `port` option to specify a TCP port number other than the default port of 22. You can only specify ports 22 and 1024 through 65535.

Using ACLs to Control SSH Access. You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you configure an ACL policy to permit or deny a specific list of IP

addresses and subnet masks for the SSH port. You must create an ACL policy file before you can use the `access-profile` option. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to cancel a previously configured ACL.

Creating an ACL Policy File. To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Chapter 11, “Policy Manager,”](#) and [Chapter 12, “Access Lists \(ACLs\),”](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile_2.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing SSH Information. To view the status of SSH2 sessions on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions and whether a valid key is present.

Example

The following command enables the SSH2 feature:

```
enable ssh2
```

The next example assumes you have already created an ACL to apply to SSH.

The following command applies the ACL `MyAccessProfile_2` to SSH:

```
enable ssh2 access-profile MyAccessProfile_2
```

History

This command was first available in the ExtremeXOS 11.0 SSH module.

The `access-profile` and `none` options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable tacacs

```
enable tacacs
```

Description

Enables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After they have been enabled, all web and Telnet logins are sent to one of the two TACACS+ servers for login name authentication.

Example

The following command enables TACACS+ user authentication:

```
enable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable tacacs-accounting

```
enable tacacs-accounting
```

Description

Enables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If accounting is used, the TACACS+ client must also be enabled.

Example

The following command enables TACACS+ accounting for the switch:

```
enable tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable tacacs-authorization

```
enable tacacs-authorization
```

Description

Enables CLI command authorization.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed. TACACS+ authentication must also be enabled to use TACACS+ authorization. Use the following command to enable authentication:

```
enable tacacs
```

Example

The following command enables TACACS+ command authorization for the switch:

```
enable tacacs-authorization
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable web http

```
enable web http
```

Description

Enables hypertext transfer protocol (HTTP) access to the switch on the default HTTP port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If HTTP access has been disabled, use this command to enable HTTP access to the switch.

Example

The following command enables HTTP on the default port:

```
enable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable web https

```
enable web https
```

Description

Enables secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to allow users to connect using a more secure HTTPS connection.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, “Software Upgrade and Boot Options,”](#) of the *ExtremeXOS Concepts Guide*.

To use secure HTTP access (HTTPS) for web-based login connections, you must specify HTTPS as the protocol when configuring the redirect URL. For more information about configuring the redirect URL, see the [configure netlogin redirect-page](#) command.

Prior to ExtremeXOS 11.2, the SSH module did not include SSL. To use SSL, you must download and install the current SSH module.

If you are currently running SSH with ExtremeXOS 11.0 or 11.1, and you want to use SSL for secure HTTPS web-based login, you must upgrade your core software image to ExtremeXOS 11.2 or later, install the SSH module that works in concert with that core software image, and reboot the switch.

Example

The following command enables SSL on the default port:

```
enable web https
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

scp2

```
scp2 {cipher [3des | blowfish]} {port <portnum>} <user>@ [<hostname> |
<ipaddress>]:<remote_file> <local_file> {vr <vr_name>}
```

or

```
scp2 {cipher [3des | blowfish]} {port <portnum>} <local_file> <user>@
[<hostname> | <ipaddress>]:<remote_file> {vr <vr_name>}
```

Description

The first command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the remote system to the switch.

The second command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the switch to a remote system.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
user	Specifies a login name for the remote host.
hostname	Specifies the name of the remote host.
ipaddress	Specifies the IP address of the remote host.
remote_file	Specifies the name of the remote file (configuration file, policy file, image file, public key file) to be transferred.
local_file	Specifies the name of the local file (configuration file, policy file, image file, public key file) to be transferred.
vr_name	Specifies the virtual router. The default virtual router is VR-Mgmt. NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

You must be running the SSH2 module (ssh.xmod), which is under Export Control, in order to use the SCP2 command.

SSH2 does not need to be enabled on the switch in order to use this command.

This command logs into the remote host as <user> and accesses the file <remote_file>. You will be prompted for a password from the remote host, if required.

Host Name, User Name, and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:)
- At symbol (@) Permitted only for user names
- Slash (/) Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command copies the configuration file test.cfg on host system1 to the switch:

```
scp2 admin@system1:test.cfg localtest.cfg
```

The following command copies the configuration file engineering.cfg from the switch to host system1:

```
scp2 engineering.cfg admin@system1:engineering.cfg
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show dhcp-server

```
show dhcp-server {vlan <vlan_name>}
```

Description

Displays the DHCP server's configuration and address allocation on a specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.

Default

N/A.

Usage Guidelines

If no VLAN is specified, the configuration and address allocation for the servers on all the VLANs is displayed.

Example

The following command displays the configuration and address allocation for the DHCP server for the VLAN *corp*:

```
show dhcp-server vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show dos-protect

```
show dos-protect {detail}
```

Description

Displays DoS protection configuration and state.

Syntax Description

detail	Specifies to display statistics in addition to configuration and state.

Default

N/A.

Usage Guidelines

Use this command to display the DoS protection settings. Using the `detail` option will also display the following cumulative statistics:

- trusted
- notify
- alerts

Example

The following command displays the DoS protection settings for the switch:

```
show dos-protect
```

The following is sample output from this command:

```
dos-protect is disabled
```

```
dos-protect settings:
```

```
interval:      1  (measurement interval secs)
```

```
acl expire time: 5  (secs)
```

```
trusted ports:
```

```
no trusted ports configured
```

```
type L3-Protect:
```

```
notify threshold: 3500  (level to log a message)
```

```
alert threshold: 4000  (level to generate an ACL)
```

The following command displays detailed DoS protection settings for the switch:

```
show dos-protect detail
```

The following is sample output from this command:

```
dos-protect is enabled
```

```
dos-protect settings:
```

```
interval:      1  (measurement interval secs)
```

```
acl expire time: 5  (secs)
```

```
trusted ports:
```

```
1:2
```

```
type L3-Protect:
```

```
  notify threshold: 3500  (level to log a message)
```

```
  alert threshold:  4000  (level to generate an ACL)
```

```
dos-protect statistics:
```

```
  trusted: 1301
```

```
  notify:  0
```

```
  alerts:  0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show ip-security anomaly-protection notify cache ports

```
show ip-security anomaly-protection notify cache ports <port list>
```

Description

Displays most anomaly notification caches.

Syntax Description

port list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

This command displays most anomaly notification caches.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X450a, X450e, and X250e platforms, whether or not included in a SummitStack, and the Aspen platform with any of the G48Te, G48Pe, G48Ta, G48Xa, 10G4Ca, or 10G4Xa blades.

show ip-security arp gratuitous-protection

```
show ip-security arp gratuitous-protection
```

Description

If configured for gratuitous ARP, displays the gratuitous ARP protection configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the name of each VLAN configured for gratuitous ARP.

If you do not have gratuitous ARP configured, the switch does not display any VLAN information.

Example

The following command displays the gratuitous ARP configuration on the switch:

```
show ip-security arp gratuitous-protectection
```

The following is sample output from this command:

```
Gratuitous ARP Protection enabled on following VLANs:  
Default, test
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security arp learning vlan

```
show ip-security arp learning {vlan} <vlan_name>
```

Description

Displays how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN.

Default

N/A.

Usage Guidelines

The switch displays the following ARP learning information:

- Port—The member port of the VLAN.
- Learn-from—The method the port uses to build the ARP table. The methods are:
 - ARP—ARP learning is enabled. The switch uses a series of requests and replies to build the ARP table.
 - DHCP—DHCP secured ARP is enabled. The switch uses DHCP snooping to build the ARP table.
 - None—Both DHCP secured ARP and ARP learning are disabled.

Example

The following command displays how the switch builds its ARP table for the VLAN learn:

```
show ip-security arp learning vlan learn
```

The following is sample output from this command:

```
Port          Learn-from
-----
2:1           ARP
2:2           DHCP,  poll 300 sec, retries 3
2:3           ARP
2:4           None
2:5           ARP
2:6           ARP
2:7           ARP
2:8           ARP
```


History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security arp validation vlan

```
show ip-security arp validation {vlan} <vlan_name>
```

Description

Displays ARP validation information for the specified VLAN.

Syntax Description

vlan_name	Specifies the name of the VLAN.

Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation-action—Determines what action(s) the switch takes when an invalid ARP is received.

Example

The following command displays ARP validation on for the VLAN valid:

```
show ip-security arp validation vlan valid
```

The following is sample output from this command:

```
-----
Port      Validation      Violation-action
-----
7         DHCP             drop-packet, block-port for 120 seconds, snmp-trap
23        DHCP             drop-packet, block-port for 120 seconds, snmp-trap
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security dhcp-snooping entries vlan

```
show ip-security dhcp-snooping entries {vlan} <vlan_name>
```

Description

Displays the DHCP bindings database on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP bindings database information:

- VLAN—The name of the DHCP-snooping VLAN
- IP Addr—The IP address of the untrusted interface or client
- MAC Addr—The MAC address of the untrusted interface or client
- Port—The port number where the untrusted interface or client attempted to access the network

Example

The following command displays the DHCP bindings database on the switch:

```
show ip-security dhcp-snooping entries vlan dhcpVlan
```

The following is sample output from this command:

```
-----
Vlan: dhcpVlan
-----
IP Addr          MAC Addr          Server  Client
-----          -
172.16.100.9     00:90:27:c6:b7:65  1:1     1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security dhcp-snooping violations vlan

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Description

Displays the MAC addressed from which the rouge DHCP packet was received by the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- Port—The specific port that received the rouge DHCP packet
- Violating MAC—The MAC address from which the rouge DHCP was received by the switch

Example

The following command displays the DHCP snooping violations for the VLAN green:

```
show ip-security dhcp-snooping violations green
```

The following is sample output from this command:

```
Violations seen on following ports
-----
Port          Violating MAC
-----
2:3           00-0c-11-a0-3e-12
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security dhcp-snooping vlan

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

Description

Displays the DHCP snooping configurations on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- DHCP snooping enabled on ports—The ports that have DHCP snooping enabled
- Trusted ports—The ports configured as trusted ports
- Trusted DHCP servers—The servers configured as trusted DHCP servers
- Port—The specific port that has DHCP snooping enabled
- Violation-action—The action the switch takes upon detecting a rogue DHCP packet on the port

Example

The following command displays the DHCP snooping settings for the switch:

```
show ip-security dhcp-snooping vlan dhcpVlan
```

The following is sample output from this command:

```
DHCP Snooping enabled on ports: 1:2, 1:3, 1:4, 1:7, 1:9
Trusted Ports: 1:7
Trusted DHCP Servers: None
-----
Port      Violation-action
-----
1:2       none
1:3       drop-packet
1:4       drop-packet, block-mac permanently
1:7       none
1:9       drop-packet, snmp-trap
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ip-security source-ip-lockdown

```
show ip-security source-ip-lockdown
```

Description

Displays the source IP lockdown configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the following source IP lockdown information:

- Port—Indicates the port that has DHCP snooping enabled and is configured for source IP lockdown
- Locked IP Address—Indicates a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address

Example

The following command displays the source IP configuration on the switch:

```
show ip-security source-ip-lockdown
```

The following is sample output from this command:

Ports	Locked IP Address
23	10.0.0.101

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show mac-lockdown-timeout fdb ports

```
show mac-lockdown-timeout fdb ports [all | <port_list>]
```

Description

Displays the MAC entries that are learned on the specified port or group of ports or for all ports on the switch along with the aging time of each port.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

If a port is down, the command displays all of the MAC entries that are maintained locally in the software.

The MAC entries learned on the specified port are displayed only if the MAC lock down timeout feature is enabled on the port. If you specify a port on which this feature is disabled, the MAC entries learned on that port are not displayed.

The switch displays the following information:

- Mac—The MAC address that defines the entry
- Vlan—The VLAN name and ID for the entry
- Age—The age of the entry, in seconds
- Flags—Flags that define the type of entry:
 - B—Egress Blackhole
 - b—Ingress Blackhole
 - F—Entry in the hardware FDB
 - L—Entry in the software
- Port—The port on which the MAC address has been learned

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3 and 2:4:

```
show mac-lockdown-timeout fdb ports 2:3, 2:4
```


The following is sample output from this command:

```

Mac              Vlan      Age  Flags  Port
-----
00:00:01:02:03:04  v1(4094)  0010 F      2:3
00:00:01:00:00:02  v1(4094)  0030 FB b   2:3
00:00:0A:02:03:04  v2(4093)  0050 L      2:4
00:00:0B:02:03:04  v2(4093)  0090 F      2:4
Flags : (F) Entry as in h/w FDB, (L) Entry in s/w and not in h/w
        (B) Egress Blackhole, (b) Ingress Blackhole

Total: 4  Entries in FDB: 3  Entries in s/w: 1

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show mac-lockdown-timeout ports

```
show mac-lockdown-timeout ports [all | <port_list>]
```

Description

Displays information about the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The switch displays the following MAC address timeout information:

- Port—Indicates the port number that you specified in the command
- MAC Lockdown Timeout—Specifies the enabled/disabled state of the MAC address lock down timeout feature.
- Timeout (in seconds)—Specifies the timeout value for the specified ports. By default, the timeout value is 15 seconds. Even if MAC address lock down is disabled, the default timeout value is displayed.

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3, 2:4, and 2:6:

```
show mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

The following is sample output from this command:

```
Ports      MAC Lockdown Timeout    Timeout (in seconds)
=====
2:3                Enabled                300
2:4                Enabled                300
2:6                Disabled               15
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show radius

```
show radius {mgmt-access | netlogin}
```

Description

Displays the current RADIUS client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS authentication server.
netlogin	Specifies configuration and statistics for the network login RADIUS authentication server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays the status of RADIUS and RADIUS accounting (enabled or disabled) and the primary and secondary servers for RADIUS and RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS configuration details related to management access.

Use the `netlogin` keyword to only RADIUS configuration details related to network login.

Example

The following command displays the current RADIUS client configuration and statistics for both management and network login:

```
show radius
```

The following is sample output from this command:

```
Switch Management Radius: enabled
Switch Management Radius server connect time out: 3 seconds
Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius: enabled
Netlogin Radius server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds
```

Primary Switch Management Radius server:

```

Server name      :
IP address       : 10.100.1.100
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0              Client timeouts     : 0
Bad authenticators: 0              Unknown types       : 0
Round Trip Time  : 0

```

Secondary Switch Management Radius server:

```

Server name      :
IP address       : 10.100.1.101
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0              Client timeouts     : 0
Bad authenticators: 0              Unknown types       : 0
Round Trip Time  : 0

```

Primary Netlogin Radius server:

```

Server name      :
IP address       : 10.100.1.200
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0              Client timeouts     : 0
Bad authenticators: 0              Unknown types       : 0
Round Trip Time  : 0

```

Secondary Netlogin Radius server:

```

Server name      :
IP address       : 10.100.1.201
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0              Client timeouts     : 0
Bad authenticators: 0              Unknown types       : 0
Round Trip Time  : 0

```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show radius-accounting

```
show radius-accounting {mgmt-access | netlogin}
```

Description

Displays the current RADIUS accounting client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS accounting server.
netlogin	Specifies configuration and statistics for the network login RADIUS accounting server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays information about the status and configuration of RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS accounting configuration details related to management access.

Use the `netlogin` keyword to display only RADIUS accounting configuration details related to network login.

Example

The following command displays RADIUS accounting client configuration and statistics for both management and network login:

```
show radius-accounting
```

The following is sample output from this command:

```
Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds

Primary Switch Management Accounting server:
  Server name      :
  IP address       : 10.100.1.100
  Server IP Port   : 1813
  Client address   : 10.116.3.101 (VR-Mgmt)
  Shared secret    : g~`#uovpkpvi~`
```

```

Acct Requests      : 0           Acct Responses      : 0
Acct Retransmits   : 0           Timeouts            : 0

```

Secondary Switch Management Accounting server:

```

Server name       :
IP address        : 10.100.1.101
Server IP Port    : 1813
Client address    : 10.116.3.101 (VR-Mgmt)
Shared secret     : g~`#uovpkkpvi~`
Acct Requests     : 0           Acct Responses     : 0
Acct Retransmits  : 0           Timeouts          : 0

```

Primary Netlogin Accounting server:

```

Server name       :
IP address        : 10.100.1.200
Server IP Port    : 1813
Client address    : 10.116.3.101 (VR-Mgmt)
Shared secret     : g~`#uovpkkpvi~`
Acct Requests     : 0           Acct Responses     : 0
Acct Retransmits  : 0           Timeouts          : 0

```

Secondary Netlogin Accounting server:

```

Server name       :
IP address        : 10.100.1.201
Server IP Port    : 1813
Client address    : 10.116.3.101 (VR-Mgmt)
Shared secret     : g~`#uovpkkpvi~`
Acct Requests     : 0           Acct Responses     : 0
Acct Retransmits  : 0           Timeouts          : 0

```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show sshd2 user-key

```
show sshd2 user-key {<key_name> {users}}
```

Description

Displays the user names bound to a key.

Syntax Description

key_name	Specifies the name of the public key to be deleted.
users	Specifies the name of the users.

Default

N/A.

Usage Guidelines

This command displays the names of the users that are bound to a public key.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show ssl

```
show ssl
```

Description

Displays the secure socket layer (SSL) configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information:

- HTTPS port configured. This is the port on which the clients will connect.
- Length of the RSA key (the number of bits used to generate the private key).
- Basic information about the stored certificate.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Appendix A, “Software Upgrade and Boot Options,”](#) of the *ExtremeXOS Concepts Guide*.

If you attempt to use this command before installing the SSH module, the switch displays a message similar to the following:

```
SSL Module: Not Installed.
```



NOTE

The switch utilizes the SSH module for SSL functionality. You do not install an SSL module, only the SSH module.

Example

The following command displays the SSL configuration:

```
show ssl
```

The following is sample output from this command:

```
HTTPS Port Number: 443
Private Key matches with the Public Key in certificate. (or Private key does not match
with the Public Key in the certificate)
RSA Key Length: 1024
Certificate:
Data:
Version: 1 (0x0)
Serial Number: 6 (0x6)
Signature Algorithm: md5WithRSAEncryption
Issuer: C=AU, O=CryptSoft Pty Ltd, CN=Test CA (1024 bit)
Validity
Not Before: Oct 16 22:31:03 2000 GMT
Not After : Jan 14 22:31:03 2003 GMT
Subject: C=AU, O=CryptSoft Pty Ltd, CN=Server test cert (512 bit)
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

show tacacs

```
show tacacs
```

Description

Displays the current TACACS+ configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- TACACS+—The current state of TACACS+, enabled or disabled.
- TACACS+ Authorization—The current state of TACACS+ authorization, enabled or disabled.
- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Server—Describes information about the primary TACACS+ server, including:
 - The name of the primary TACACS+ server
 - The IP address of the primary TACACS+ server
 - The TCP port to use to contact the primary TACACS+ server
 - The IP address and VR used by the switch
 - The shared secret configured for the primary TACACS+ server
- Secondary TACACS+ Server—Contains the same type of output as the primary TACACS+ server for the secondary TACACS+ server, if configured.
- TACACS+ Acct Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ accounting server failure.
- TACACS+ Accounting Server parameters, if configured. Contains the same type of output as the TACACS+ server for the TACACS+ accounting server(s), if configured.

Example

The following command displays TACACS+ client configuration and statistics:

```
show tacacs
```

The following is sample output from this command:

```
TACACS+: enabled
TACACS+ Authorization: enabled
TACACS+ Accounting : enabled
TACACS+ Server Connect Timeout sec: 3
Primary TACACS+ Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret  : qijxou
Secondary TACACS+ Server:
  Server name      :
  IP address       : 10.201.31.235
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret  : qijxou
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret  : qijxou
Secondary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.235
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret  : qijxou
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show tacacs-accounting

```
show tacacs-accounting
```

Description

Displays the current TACACS+ accounting client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Accounting Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Accounting Server—Describes information about the primary TACACS+ accounting server, including:
 - The name of the primary TACACS+ accounting server
 - The IP address of the primary TACACS+ accounting server
 - The TCP port to use to contact the primary TACACS+ accounting server
 - The IP address and VR used by the switch
 - The shared secret configured for the primary TACACS+ accounting server
- Secondary TACACS+ Accounting Server—Contains the same type of output as the primary TACACS+ accounting server for the secondary TACACS+ accounting server, if configured.

Example

The following command displays TACACS+ accounting client configuration and statistics:

```
show tacacs-accounting
The following is sample output of this command:
TACACS+ Accounting : enabled
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port   : 49
  Client address   : 10.201.31.85 (VR-Default)
  Shared secret    : qijxou
Secondary TACACS+ Accounting Server:Not configured
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vlan dhcp-address-allocation

```
show vlan <vlan_name> dhcp-address-allocation
```

Description

Displays the DHCP server's address allocation on a specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP for the VLAN *corp*:

```
show vlan corp dhcp-address-allocation
```

The following is sample output from this command:

```
=====
IP                MAC                State      Lease Renewal Time
=====
      10.0.0.2     00:02:03:04:05:00   Offered    0000:00:10
      10.0.0.3     00:08:03:04:05:00   Assigned    0000:59:09
      10.0.0.4     ee:1c:00:04:05:00   Assigned    0000:59:09
=====
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show vlan dhcp-config

```
show vlan <vlan_name> dhcp-config
```

Description

Displays the DHCP server's configuration for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP server for the VLAN *corp*:

```
show vlan corp dhcp-config
```

The following is sample output from this command:

```
DHCP Address Range   : 10.0.0.2->10.0.0.20
DHCP Lease Timer     : Not configured (Default = 7200 seconds)
Default Gateway      : 10.0.0.1
DNS Server           : 10.0.0.80
WINS Server          : 10.0.0.80
Ports DHCP Enabled   : 1:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show vlan security

```
show vlan <vlan_name> security
```

Description

Displays the MAC limit-learning and lock-learning information for the specified VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

The switch displays the following information:

- Port—Indicates the port on which the MAC address has been learned
- Limit—Indicates that there is either a limited or unlimited amount of learned entries
- State—Indicates that the current FDB entries for the port are permanent, no additional entries are learned, or that the port allows unlimited, dynamic learning
- Learned—Specifies the number of learned entries
- Blackholed—Specifies the number of blackholed entries
- Locked—Specifies the number of locked entries

Example

The following command displays the security setting of the DHCP server for the VLAN *corp*:

```
show vlan blue security
```

The following is sample output from this command:

Port	Limit	State	Learned	Blackholed	Locked
24	Unlimited	Unlocked	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

ssh2

```
ssh2 {cipher [3des | blowfish]} {port <portnum>} {compression [on | off]}
{user <username>} {<username>@} [<host> | <ipaddress>] {<remote command>}
{vr <vr_name>}
```

Description

Initiates an SSH2 client session to a remote SSH2 server.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
on	Specifies that the data is to be compressed.
off	Specifies that compression is not to be used. This is the default.
username	Specifies a login name for the remote host, as an alternate to the <code>username@host</code> parameter. Can be omitted if it is the same as the username on the switch.
host	Specifies the name of the remote host.
ipaddress	Specifies the IP address of the remote host.
remote command	Specifies a command to be passed to the remote system for execution. The switch does not support remote commands. The option is only valid if the remote system is a system, such as a UNIX workstation, that accepts remote commands.
vr_name	Specifies the virtual router. The default virtual router is VR-Mgmt. NOTE: The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines



NOTE

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches do not support user-created VRs.

You must be running the SSH2 module (ssh.xmod), which is under Export Control, in order to use the SSH2 client command.

SSH2 does not need to be enabled on the switch in order to use this command.

Typically, this command is used to establish a secure session to a remote switch. You are prompted for your password. Once you have logged in successfully, all ExtremeXOS command you enter are executed on the remote switch. When you terminate the remote session, commands will then resume being executed on the original switch.

Host Name, User Name, and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:) Permitted for host names and remote IP addresses
- At symbol (@) Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command establishes an SSH2 session on switch engineering1:

```
ssh2 admin@engineering1
```

The following command establishes and SSH2 session with the switch named BlackDiamond8810 over TCP port 2050 with compression enabled:

```
ssh2 compression on port 2050 admin@BlackDiamond8810
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms with the SSH2 module installed.

unconfigure radius

```
unconfigure radius {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS client configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Unconfigures the primary RADIUS server.
secondary	Unconfigures the secondary RADIUS server.

Default

Unconfigures both primary and secondary servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary servers for management and network login.

The following list describes the available keywords:

- **mgmt-access**—Use this keyword to unconfigure only the server(s) for management functions.
- **netlogin**—Use this keyword to unconfigure only the server(s) for network login.
- **primary**—Use this keyword to specify only the primary RADIUS sever.
- **secondary**—Use this keyword to specify only the secondary RADIUS server.

Example

The following command unconfigures the secondary RADIUS server settings for both management and network login:

```
unconfigure radius server secondary
```

The following command unconfigures the secondary RADIUS server settings for only network login:

```
unconfigure radius netlogin server secondary
```

The following command unconfigures all RADIUS server settings for only management functions:

```
unconfigure radius mgmt-access
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure radius-accounting

```
unconfigure radius-accounting {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS accounting server configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
primary	Unconfigures the primary RADIUS accounting server.
secondary	Unconfigures the secondary RADIUS accounting server.

Default

Unconfigures both the primary and secondary accounting servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary accounting servers for management and network login.

The following list describes the available keywords:

- **mgmt-access**—Use this keyword to unconfigure only the accounting server(s) for management functions.
- **netlogin**—Use this keyword to unconfigure only the accounting server(s) for network login.
- **primary**—Use this keyword to specify only the primary RADIUS accounting sever.
- **secondary**—Use this keyword to specify only the secondary RADIUS accounting server.

Example

The following command unconfigures the secondary RADIUS accounting server settings for both management and network login:

```
unconfigure radius-accounting server secondary
```

The following command unconfigures the secondary RADIUS accounting server settings for only network login:

```
unconfigure radius-accounting netlogin server secondary
```

The following command unconfigures all RADIUS accounting server settings for only management functions:

```
unconfigure radius-accounting mgmt-access
```


History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure tacacs

```
unconfigure tacacs {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ server configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ server.
secondary	Unconfigures the secondary TACACS+ server.

Default

Unconfigures both the primary and secondary TACACS+ servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ servers settings:

```
unconfigure tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure tacacs-accounting

```
unconfigure tacacs-accounting {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ accounting server configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ accounting server.
secondary	Unconfigures the secondary TACACS+ accounting server.

Default

Unconfigures both the primary and secondary TACACS+ accounting servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ accounting servers settings:

```
unconfigure tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure trusted-ports trust-for dhcp-server

```
unconfigure trusted-ports [<ports>|all] trust-for dhcp-server
```

Description

Unconfigures, disables one or more DHCP trusted ports.

Syntax Description

ports	Specifies one or more trusted ports.
all	Specifies all trusted ports.

Default

N/A.

Usage Guidelines

Use this command to disable one or more DHCP trusted ports.

Displaying DHCP Trusted Server Information. To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command unconfigures ports 2:2 and 2:3 as trusted ports:

```
unconfigure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

unconfigure vlan dhcp

```
unconfigure vlan <vlan_name> dhcp
```

Description

Unconfigure all the DHCP configuration information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP server for the VLAN *temporary*:

```
unconfigure temporary dhcp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure vlan dhcp-address-range

```
unconfigure vlan <vlan_name> dhcp-address-range
```

Description

Unconfigure the DHCP address range information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP address range for the VLAN *temporary*:

```
unconfigure temporary dhcp-address-range
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure vlan dhcp-options

```
unconfigure vlan <vlan_name> dhcp-options
```

Description

Unconfigure the DHCP option information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP options for the VLAN *temporary*:

```
unconfigure temporary dhcp-options
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

20 Network Login Commands

This chapter describes commands for configuring network login.

Network login is a feature designed to control the admission of user packets into a network by giving network access only to users that have been properly authenticated. Network login is controlled by an administrator on a per port, per VLAN basis and uses an integration of DHCP, user authentication over the web interface, user authentication by MAC address, or 802.1x client software, and a RADIUS server to provide a user database or specific configuration details.

Network login has two modes of operation:

- Campus mode, used when a port in a VLAN will move to another VLAN when authentication has been completed successfully. This mode is for the roaming user who will not always be using the same port for authentication. Campus mode requires a DHCP server and a RADIUS server configured for Extreme Network Login.
- ISP mode, used when the port and VLAN used will remain constant. All network settings are configured for that VLAN.

A DHCP server is included to support Network Login functionality.

clear netlogin state

```
clear netlogin state [{port <portlist>} {vlan <vlan name>}]
```

Description

Clears and initializes the Network Login sessions on a VLAN port.

Syntax Description

portlist	Specifies the ports to clear.
vlan name	Specifies a VLAN to clear.

Default

None.

Usage Guidelines

Clear the states of every MAC learned on this VLAN port and put the port back to unauthenticated state. The port will be moved to its original VLAN if configured in Campus mode.

Example

The following command clears the Network Login state of port 2:9 in VLAN *corp*:

```
clear netlogin state port 2:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

clear netlogin state mac-address

```
clear netlogin state mac-address <mac>
```

Description

Initialize/Reset the Network Login sessions for a specified supplicant.

Syntax Description

mac	Specifies the MAC address of the supplicant.

Default

N/A.

Usage Guidelines

This command is essentially equivalent to a particular supplicant logging out. The MAC address will be cleared from the FDB, the port is put back to its original VLAN (for Campus mode), and the port state is set to unauthenticated, if this was the last authenticated MAC on this port.

Example

The following command resets the Network Login session for the supplicant with the MAC address of 00:e0:18:01:32:1f:

```
clear netlogin state mac-address 00:e0:18:01:32:1f
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin add mac-list

```
configure netlogin add mac-list [<mac> {<mask>} | default] {encrypted}
{<password>} {ports <port_list>}
```

Description

Adds an entry to the MAC address list for MAC-based Network Login.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

If no password is specified, the MAC address will be used.

Usage Guidelines

Use this command to add an entry to the MAC address list used for MAC-based Network Login.

If no match is found in the table of MAC entries, and a default entry exists, the default will be used to authenticate the client. All entries in the list are automatically sorted in longest prefix order.

Associating a MAC Address to a Port. You can configure the switch to accept and authenticate a client with a specific MAC address. Only MAC addresses that have a match for the specific ports are sent for authentication. For example, if you associate a MAC address with one or more ports, only authentication requests for that MAC addresses received on the port(s) are sent to the RADIUS server. The port(s) block all other authentication requests that do not have a matching entry. This is also known as secure MAC.

To associate a MAC address with one or more ports, specify the `ports` option when using the `configure netlogin add mac-list [<mac> {<mask>} | default] {encrypted} {<password>} {ports <port_list>}` command.

You must enable MAC-based netlogin on the switch and the specified ports before using this command. If MAC-based netlogin is not enabled on the specified port(s), the switch displays a warning message similar to the following:

```
WARNING: Not all specified ports have MAC-Based NetLogin enabled.
```

If this occurs, make sure to enable MAC-based netlogin.

Example

The following command adds the MAC address 10:20:30:40:50:60 with the password *foo* to the list:

```
configure netlogin add mac-list 10:20:30:40:50:60 password foo
```

The following command associates MAC address 10:20:30:40:50:70 with ports 2:2 and 2:3. This means authentication requests from MAC address 10:20:30:40:50:70 are only accepted on ports 2:2 and 2:3:

```
configure netlogin add mac-list mac 10:20:30:40:50:70 ports 2:2-2:3
```

History

This command was first available in ExtremeXOS 11.1.

The `ports` option was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure netlogin agingtime

```
configure netlogin agingtime <minutes>
```

Description

Lets you configure netlogin aging.

Syntax Description

minutes	Specifies the aging time in minutes.

Default

The default value is 5.

Usage Guidelines

Use this command to configure the aging time for netlogin. The aging time is the time after which learned clients that failed authentication or did not attempt to authenticate are removed from the system. This prevents the switch from keeping all clients ever seen on a netlogin-enabled port.

The range can be from 0 to 3000, where 0 indicates no age out.

Example

The following command specifies an aging time of 15 minutes:

```
configure netlogin agingtime 15
```

History

This command was first available in ExtremeXOS 11.1.

configure netlogin authentication database-order

```
configure netlogin [mac | web-based] authentication database-order  
[[radius] | [local] | [radius local] | [local radius]]
```

Description

Configures the order of database authentication protocols to use.

Syntax Description

mac	Specifies mac-based authentication.
web-based	Specifies Web-based authentication.
radius	Specifies an authentication order from only the Radius database.
local	Specifies an authentication order from only the local database.
radius local	Specifies an authentication order of Radius database first, followed by local.
local radius	Specifies an authentication order of local database first, followed by Radius.

Default

By default, the authentication order is Radius, Local-User database.

Usage Guidelines

Use this command in situations where, when both a NetLogin Radius server and a Local-User database are configured, you want to have control over which database to use first. If one authentication fails, the other database is tried; if that authentication is successful, the switch authenticates the NetLogin user.

Example

The following command sets the database authentication order to Local-User database, Radius:

```
configure netlogin mac authentication database-order local radius
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure netlogin banner

```
configure netlogin banner <banner>
```

Description

Configures the Network Login page banner.

Syntax Description

banner	Specifies the HTML code for the banner.

Default

The default banner is the Extreme Networks logo.

Usage Guidelines

The banner is a quoted, HTML string, that will be displayed on the Network Login page. The string is limited to 1024 characters.

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the netlogin page banner:

```
configure netlogin banner "<html><head>Please Login</head></html>"
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin base-url

```
configure netlogin base-url <url>
```

Description

Configures the base URL for Network Login.

Syntax Description

url	Specifies the base URL for Network Login.

Default

The base URL default value is "network-access.com."

Usage Guidelines

When you login using a web browser, you are redirected to the specified base URL, which is the DNS name for the switch.

You must configure a DNS name of the type "www.xx...xx.xxx" or "xx...xx.xxx".

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the netlogin base URL as access.net:

```
configure netlogin base-url access.net
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin delete mac-list

```
configure netlogin delete mac-list [<mac> {<mask>} | default]
```

Description

Deletes an entry from the MAC address list for MAC-based Network Login.

Syntax Description

mac	Specifies the MAC address to delete.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the MAC address list used for MAC-based Network Login.

Example

The following command deletes the MAC address 10:20:30:40:50:60 from the list:

```
configure netlogin delete mac-list 10:20:30:40:50:60
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x eapol-transmit-version

```
configure netlogin dot1x eapol-transmit-version <eapol-version>
```

Description

Configures the default EAPOL version sent in transmitted packets for Network Login.

Syntax Description

eapol-version	Specifies the EAPOL version. Choices are "v1" or "v2".

Default

The default is "v1".

Usage Guidelines

Although ExtremeXOS supports EAPOL version 2, some clients do not yet accept the version 2 EAPOL packets. The packet format for the two versions is the same.

Example

The following command changes the EAPOL version to 2:

```
configure netlogin dot1x eapol-transmit-version v2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x guest-vlan

```
configure netlogin dot1x guest-vlan <vlan_name> {ports <port_list>}
```

Description

Configures a guest VLAN for 802.1x authentication netlogin.

Syntax Description

vlan_name	Specifies the name of the guest VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

N/A.

Usage Guidelines

This command configures the guest VLAN for 802.1x on the current virtual router (VR).



NOTE

Beginning with ExtremeXOS 11.6, you can configure guest VLANs on a per port basis, which allows you to configure more than one guest VLAN per VR. In ExtremeXOS 11.5 and earlier, you can only configure guest VLANs on a per VLAN basis, which allows you to configure only one guest VLAN per VR.

If you do not specify any ports, the guest VLAN is configured for all ports.

Each port can have a different guest VLAN.

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Keep in mind the following when configuring guest VLANs:

- You must create a VLAN and configure it as a guest VLAN before enabling the guest VLAN feature.
- Configure guest VLANs only on netlogin ports with 802.1x enabled.
- Movement to guest VLANs is not supported on netlogin ports with MAC-based or web-based authentication.
- 802.1x must be the only authentication method enabled on the port for movement to guest VLAN.
- No supplicant on the port has 802.1x capability.
- You configure only one guest VLAN per virtual router interface.



NOTE

The supplicant does not move to a guest VLAN if it fails authentication after an 802.1x exchange; the supplicant moves to the guest VLAN only if it does not respond to an 802.1x authentication request.

Modifying the Supplicant Timer. By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is not a user-configured parameter.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-period
<quiet_period>} {reauth-period <reauth_period>} {supp-resp-timeout
<supp_resp_timeout>}]
```

If a supplicant on a port in the guest VLAN becomes 802.1x-capable, the switch starts processing the 802.1x responses from the supplicant. If the supplicant is successfully authenticated, the port moves from the guest VLAN to the destination VLAN specified by the RADIUS server.

Enabling Guest VLANs. To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Example

The following command creates a guest VLAN for 802.1x named guest for all ports:

```
configure netlogin dot1x guest-vlan guest
```

The following command creates a guest VLAN named guest for ports 2 and 3:

```
configure netlogin dot1x guest-vlan guest ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

The `ports` option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x timers

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-
period <quiet_period>} {reauth-period <reauth_period>} {supp-resp-timeout
<supp_resp_timeout>}]
```

Description

Configures the 802.1x timers for Network Login.

Syntax Description

server-timeout	Specifies the timeout period for a response from the RADIUS server. The range is 1 to 120 seconds.
quiet-period	Specifies the time for which the switch will not attempt to communicate with the supplicant after authentication has failed. The range is 0 to 65535 seconds.
reauth-period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 7200 seconds.
supp-resp-timeout	Specifies the time for which the switch will wait for a response from the supplicant. The range is 1 to 120 seconds.

Default

The defaults are as follows:

- server-timeout—30 seconds.
- quiet-period—60 seconds.
- reauth-period—3600 seconds.
- supp-resp-timeout—30 seconds.

Usage Guidelines

To disable re-authentication, specify 0 for the `reauth-period` parameter.

If you attempt to configure a timer value that is out of range (not supported), the switch displays an error message. The following is a list of sample error messages:

- server-timeout—ERROR: RADIUS server response timeout out of range (1..120 sec)
- quiet-period—%% Invalid number detected at '^' marker.
%% Input number must be in the range [0, 65535].
- reauth-period—ERROR: Re-authentication period out of range (0, 30..7200 sec)
- supp-resp-timeout—ERROR: Supplicant response timeout out of range (1..120 sec)

To display the 802.1x timer settings, use the `show netlogin` and `show netlogin dot1x` commands.

Example

The following command changes the 802.1x server-timeout to 10 seconds:

```
configure netlogin dot1x timers server-timeout 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin dynamic-vlan

```
configure netlogin dynamic-vlan [disable | enable]
```

Description

Configures the switch to automatically and dynamically create a VLAN after receiving authentication requests from one or more supplicants (clients).

Syntax Description

disable	Specifies that the switch does not automatically create dynamic VLANs. This is the default behavior.
enable	Specifies that the switch automatically create dynamic VLANs.

Default

The default is disabled.

Usage Guidelines

Use this command to configure the switch to dynamically create a VLAN. If configured for dynamic VLAN creation, the switch automatically creates a supplicant VLAN that contains both the supplicant's physical port and one or more uplink ports.

A dynamically created VLAN is only a Layer 2 bridging mechanism; this VLAN does not work with routing protocols to forward traffic. After the switch unauthenticates all of the supplicants from the dynamically created VLAN, the switch deletes that VLAN.



NOTE

Dynamically created VLANs do not support the session refresh feature of web-based netlogin because dynamically created VLANs do not have an IP address.

By dynamically creating and deleting VLANs, you minimize the number of active VLANs configured on your edge switches. In addition, the RADIUS server forwards VSA information to dynamically create the VLAN thereby simplifying switch management. A key difference between dynamically created VLANs and other VLANs is that the switch does not save dynamically created VLANs. Even if you use the `save` command, the switch does not save a dynamically created VLAN.

Supported Vendor Specific Attributes. To prevent conflicts with existing VLANs on the switch, the RADIUS server uses Vendor Specific Attributes (VSAs) to forward VLAN information, including VLAN ID, to the switch. The following list specifies the supported VSAs for configuring dynamic netlogin VLANs:

- Extreme: Netlogin-VLAN-ID (VSA 209)
- Extreme: Netlogin-Extended-VLAN (VSA 211)

**NOTE**

If the ASCII string only contains numbers, it is interpreted as the VLAN ID. Dynamic VLANs only support numerical VLAN IDs; VLAN names are not supported.

- IETF: Tunnel-Private-Group-ID (VSA 81)

The switch automatically generates the VLAN name in the following format: `SYS_NLD_<TAG>` where `<TAG>` specifies the VLAN ID. For example, a dynamic netlogin VLAN with an ID of 10 has the name `SYS_NLD_0010`. >

Specifying the Uplink Ports. To specify one or more ports as tagged uplink ports that are added to the dynamically created VLAN, use the following command:

```
configure netlogin dynamic-vlan uplink-ports [<port_list> | none]
```

The uplink ports send traffic to and from the supplicants from the core of the network.

By default the setting is none. For more information about this command, see the usage guidelines for `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information. To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with `SYS_NLD_` and the output contains a `d` flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command automatically adds ports 1:1-1:2 to the dynamically created VLAN as uplink ports:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin dynamic-vlan uplink-ports

```
configure netlogin dynamic-vlan uplink-ports [<port_list> | none]
```

Description

Specifies which port(s) are added as tagged, uplink ports to the dynamically created VLANs for netlogin.

Syntax Description

port_list	Specifies one or more ports to add to the dynamically created VLAN for netlogin.
none	Specifies that no ports are added. This is the default setting.

Default

The default setting is none.

Usage Guidelines

Use this command to specify which port(s) are used as uplink ports and added to the dynamically created VLAN for netlogin. The uplink ports send traffic to and from the supplicants from the core of the network.

Uplink ports should not be configured for netlogin (netlogin is disabled on uplink ports). If you specify an uplink port with netlogin enabled, the configuration fails and the switch displays an error message similar to the following:

```
ERROR: The following ports have NetLogin enabled: 1, 2
```

If this occurs, select a port with netlogin disabled.

Enabling Dynamic Netlogin VLANs. To configure the switch to dynamically create a VLAN upon receiving an authentication response, use the following command:

```
configure netlogin dynamic-vlan [disable | enable]
```

By default, the setting is disabled. For more detailed information about this command, see the usage guidelines [configure netlogin dynamic-vlan uplink-ports](#).

Viewing Status Information. To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with `SYS_NLD_` and the output contains a `d` flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command configures the switch to add ports 1:1-1:2 to the dynamically created netlogin VLAN:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin local-user

```
configure netlogin local-user <user-name> {vlan-vsa [{tagged | untagged}
[<vlan_name>] | <vlan_tag>]] | none}}
```

Description

Configures an existing local netlogin account.

Syntax Description

user-name	Specifies the name of an existing local netlogin account.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
vlan_name	Specifies the name of the destination VLAN.
vlan_tag	Specifies the VLAN ID, tag, of the destination VLAN.
none	Specifies that the VSA 211 wildcard (*) is applied, only if you do not specify tagged or untagged

Default

N/A.

Usage Guidelines

Use this command to modify the attributes of an existing local netlogin account. You can update the following attributes associated with a local netlogin account:

- Password of the local netlogin account
- Destination VLAN attributes including: adding clients tagged or untagged, the name of the VLAN, and the VLAN ID



NOTE

Passwords are case-sensitive and must have a minimum of 1 character and a maximum of 32 characters.

You must create a local netlogin account before using this command. To create a local netlogin user name and password, use the following command:

```
create netlogin local-user <user-name> {encrypted <password>} {vlan-vsa [{tagged |
untagged} [<vlan_name>] | <vlan_tag>]]}
```

If the switch displays a message similar to the following:

```
* SummitX450.14 # configure netlogin local-user purplenet
^
%% Invalid input detected at '^' marker.
```

You might be attempting to modify a local netlogin account that is not present or the switch, or you might have incorrectly entered the account name. To confirm the names of the local netlogin accounts on your switch, use the following command:

```
show netlogin local-users
```

Additional Requirements. This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

Passwords are case-sensitive. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Example

This section contains the following examples:

- Updating the password
- Modifying destination VLAN attributes

Updating the Password. The following command updates the password of an existing local netlogin account:

```
configure netlogin local-user megtest
```

After you enter the local netlogin user name, press [Enter]. The switch prompts you to enter a password; however, the switch does not display the password. At the prompt enter the new password:

```
password:
```

After you enter the new password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

Updating VLAN Attributes. You can add a destination VLAN, change the destination VLAN, or remove the destination from an existing local netlogin account. This example changes the destination VLAN for the specified local netlogin account:

```
configure netlogin local-user megtest vlan-vsa green
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure netlogin move-fail-action

```
configure netlogin move-fail-action [authenticate | deny]
```

Description

Configures the action network login takes if a VLAN move fails. This can occur if two clients attempt to move to an untagged VLAN on the same port.

Syntax Description

authenticate	Specifies that the client is authenticated.
deny	Specifies that the client is not authenticated. This is the default setting.

Default

The default setting is `deny`.

Usage Guidelines

Use this command to specify how netlogin behaves if a VLAN move fails. Netlogin can either authenticate the client on the current VLAN or deny the client.

The following describes the parameters of this command if two clients want to move to a different untagged VLAN on the same port:

- `authenticate`—Network login authenticates the first client that requests a move and moves that client to the requested VLAN. Network login authenticates the second client but does not move that client to the requested VLAN. The second client moves to the first client's authenticated VLAN.
- `deny`—Network login authenticates the first client that requests a move and moves that client. Network login does not authenticate the second client.

To view the current move-fail-action setting on the switch, use the `show netlogin` command.

Example

The following command configures netlogin to authenticate the client on the current VLAN:

```
configure netlogin move-fail-action authenticate
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin ports mode

```
configure netlogin ports [all | <port_list>] mode [mac-based-vlans | port-based-vlans]
```

Description

Configures the netlogin port's mode of operation.

Syntax Description

all	Specifies all netlogin ports.
port_list	Specifies one or more netlogin ports.
mac-based-vlans	Allows more than one untagged VLAN.
port-based-vlans	Allows only one untagged VLAN. This is the default behavior.

Default

The default setting is `port-based-vlans`.

Usage Guidelines

Use this command to configure netlogin MAC-based VLANs on a netlogin port.

If you modify the mode of operation to `mac-based-vlans` and later disable all netlogin protocols on that port, the mode of operation automatically returns to `port-based-vlans`.

When you change the netlogin port's mode of operation, the switch deletes all currently known supplicants from the port and restores all VLANs associated with that port to their original state. In addition, by selecting `mac-based-vlans`, you are unable to manually add or delete untagged VLANs from this port. Netlogin now controls these VLANs.

With netlogin MAC-based operation, every authenticated client has an additional FDB flag that indicates a translation MAC address. If the supplicant's requested VLAN does not exist on the port, the switch adds the requested VLAN.

Important Rules and Restrictions. This section summarizes the rules and restrictions for configuring netlogin MAC-based VLANs:

- If you attempt to configure the port's mode of operation before enabling netlogin, the switch displays an error message similar to the following:

```
ERROR: The following ports do not have NetLogin enabled; 1
```

To enable netlogin on the switch, use the following command to enable netlogin and to specify an authentication method (for example, 802.1x—identified as dot1x in the CLI):

```
enable netlogin dot1x
```

To enable netlogin on the ports, use the following command to enable netlogin and to specify an authentication method (for example, 802.1x—identified as dot1x in the CLI):

```
enable netlogin ports 1:1 dot1x
```

- On ExtremeXOS versions prior to 12.0 on switches other than the Summit family and the BlackDiamond 12800 series, 10 Gigabit Ethernet ports such as those on the 10G4X I/O module and the uplink ports on the Summit family of switches do not support netlogin MAC-based VLANs.

If you attempt to configure netlogin MAC-based VLANs on 10 Gigabit Ethernet ports, the switch displays an error message similar to the following:

```
ERROR: The following ports do not support the MAC-Based VLAN mode; 1, 2, 10
```

In ExtremeXOS version 12.0, BlackDiamond 12800 series, X450 "a series" and "e series", and BlackDiamond 8800 10G4Xa support MAC-based vlans on 10 Gigabit Ethernet ports. The X450 "original series" and BD8800 10G4X 10 Gigabit Ethernet ports do not support MAC-based vlans.

- You can have a maximum of 1,024 MAC addresses per I/O module or per Summit family switch.

Displaying FDB Information. To view netlogin-related FDB entries, use the following command:

```
show fdb netlogin [all | mac-based-vlans]
```

The following is sample output from the `show fdb netlogin mac-based-vlans` command:

Mac	Vlan	Age	Use	Flags	Port List
00:04:96:10:51:80	VLONE(0021)	0086	0000	n m	v 1:11
00:04:96:10:51:81	VL TWO(0051)	0100	0000	n m	v 1:11
00:04:96:10:51:91	VL TWO(0051)	0100	0000	n m	v 1:11

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole,
 b - Ingress Blackhole, v - NetLogin MAC-Based VLAN.

The flags associated with netlogin include:

- v—Indicates the FDB entry was added because the port is part of a MAC-Based virtual port/VLAN combination.
- n—Indicates the FDB entry was added by network login.

Displaying Port and VLAN Information. To view information about the VLANs that are temporarily added in MAC-based mode for netlogin, use the following command

```
show ports <port_list> information detail
```

The following is sample output from this command:

```
Port: 1
Virtual-router: VR-Default
Type: UTP
Random Early drop: Disabled
Admin state: Enabled with auto-speed sensing auto-duplex
Link State: Active, 100Mbps, full-duplex
Link Counter: Up 1 time(s)
VLAN cfg:
  Name: Default, Internal Tag = 1(MAC-Based), MAC-limit = No-limit
...<truncated output>
Egress 802.1p Replacement: Disabled
NetLogin: Enabled
NetLogin authentication mode: Mac based
NetLogin port mode: MAC based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
```



```
auto-polarity: Enabled
```

The added output displays information about the mode of operation for the netlogin port.

- `VLAN cfg`—The term MAC-based appears next to the tag number.
- `Netlogin port mode`—This output was added to display the port mode of operation. Mac based appears and the network login port mode of operation.

To view information about the ports that are temporarily added in MAC-based mode for netlogin, due to discovered MAC addresses, use the following command:

```
show vlan detail
```

The following is sample output from this command:

```
VLAN Interface with name Default created by user
Tagging:      802.1Q Tag 1
Priority:      802.1P Priority 0
Virtual router: VR-Default
STPD:         s0(Disabled,Auto-bind)
Protocol:      Match all unfiltered protocols
Loopback:      Disable
NetLogin:      Disabled
Rate Shape:    Disabled
QosProfile:    None configured
Ports: 26.      (Number of active ports=2)
  Untag:      *1um,      *2,      3,      4,      5,      6,      7,
               8,      9,      10,      11,      12,      13,      14,
               15,      16,      17,      18,      19,      20,      21,
               22,      23,      24,      25,      26
Flags: (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (a) Authenticated NetLogin Port
        (u) Unauthenticated NetLogin port, (m) Mac-Based port
```

The flags associated with netlogin include:

- `a`—Indicates an authenticated network login port.
- `u`—Indicates an unauthenticated network login port.
- `m`—Indicates that the netlogin port operates in MAC-based mode.

Example

The following command configures the netlogin ports mode of operation:

```
configure netlogin ports 1:1-1:10 mode mac-based-vlans
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure netlogin ports no-restart

```
configure netlogin ports [all | <port_list>] no-restart
```

Description

Disables the netlogin port restart feature.

Syntax Description

all	Specifies all netlogin ports.
port_list	Specifies one or more netlogin ports.

Default

The default setting is no-restart; the netlogin port restart feature is disabled.

Usage Guidelines

Use this command to disable the netlogin port restart feature on a netlogin port.

Configure netlogin port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration. To display the netlogin settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port <port_list>
```

Output from this command includes the enable/disable state for netlogin port restart.

Example

The following command disables netlogin port restart on port 1:1:

```
configure netlogin ports 1:1 no-restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin ports restart

```
configure netlogin ports [all | <port_list>] restart
```

Description

Enables the netlogin port restart feature.

Syntax Description

all	Specifies all netlogin ports.
port_list	Specifies one or more netlogin ports.

Default

The default setting is no-restart; the netlogin port restart feature is disabled.

Usage Guidelines

Use this command to enable the netlogin port restart feature on a netlogin port. This allows netlogin to restart specific netlogin-enabled ports when the last authenticated supplicant releases, regardless of the configured protocols on the port.

Configure netlogin port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration. To display the netlogin settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port <port_list>
```

Output from this command includes the enable/disable state for netlogin port restart.

Example

The following command enables netlogin port restart on port 1:1:

```
configure netlogin ports 1:1 restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin redirect-page

```
configure netlogin redirect-page <url>
```

Description

Configures the redirect URL for Network Login.

Syntax Description

url	Specifies the redirect URL for Network Login.

Default

The redirect URL default value is "http://www.extremenetworks.com".

Usage Guidelines

In ISP mode, you can configure netlogin to be redirected to a base page after successful login using this command. If a RADIUS server is used for authentication, then base page redirection configured on the RADIUS server takes priority over this configuration.

You must configure a complete URL starting with `http://` or `https://`

To support https, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. For more information about SSH2, see [Chapter 16, "Security,"](#) in the *ExtremeXOS Concepts Guide*. For information about installing the SSH module, see [Appendix A, "Software Upgrade and Boot Options,"](#) in the *ExtremeXOS Concepts Guide*.

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the redirect URL as `http://www.extremenetworks.com/support`:

```
configure netlogin redirect-page http://www.extremenetworks.com/support
```

History

This command was first available in ExtremeXOS 11.1.

Support for https was introduced in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure netlogin vlan

```
configure netlogin vlan <vlan_name>
```

Description

Configures the VLAN for Network Login.

Syntax Description

vlan	Specifies the VLAN for Network Login.

Default

N/A.

Usage Guidelines

This command will configure the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled. NetLogin can only be enabled when a VLAN is assigned (and no ports are configured for it).

By default no VLAN is assigned for network login.

Example

The following command configures the VLAN *login* as the network login VLAN:

```
configure netlogin vlan login
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure vlan netlogin-lease-timer

```
configure vlan <vlan name> netlogin-lease-timer <seconds>
```

Description

Configures the timer value returned as part of the DHCP response for clients attached to network login-enabled ports.

Syntax Description

vlan name	Specifies the VLAN to which this timer value applies.
seconds	Specifies the timer value, in seconds.

Default

10 seconds.

Usage Guidelines

The timer value is specified in seconds.

This command applies only to the web-based authentication mode of network login.

Example

The following command sets the timer value to 15 seconds for VLAN *corp*:

```
configure vlan corp netlogin-lease-timer 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

create netlogin local-user

```
create netlogin local-user <user-name> {encrypted <password>} {vlan-vsa
[[{tagged | untagged} [<vlan_name>] | <vlan_tag>]]}
```

Description

Creates a local netlogin user name and password.

Syntax Description

user-name	Specifies a new local netlogin user name. User names must have a minimum of 1 character and a maximum of 32 characters.
encrypted	The encrypted option is used by the switch to encrypt the password. Do not use this option through the command line interface (CLI).
password	Specifies a local netlogin user password. Passwords must have a minimum of 0 characters and a maximum of 32 characters.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
vlan_name	Specifies the name of the destination VLAN.
vlan_tag	Specifies the VLAN ID, tag, of the destination VLAN.

Default

N/A.

Usage Guidelines

Use this command to create a local network login account and to configure the switch to use its local database for network login authentication. This method of authentication is useful in the following situations:

- If both the primary and secondary (if configured) RADIUS servers timeout or are unable to respond to authentication requests.
- If no RADIUS servers are configured.
- If the RADIUS server used for network login authentication is disabled.

If any of the above conditions are met, the switch checks for a local user account and attempts to authenticate against that local account.

Extreme Networks recommends creating a maximum of 64 local accounts. If you need more than 64 local accounts, Extreme Networks recommends using RADIUS for authentication. For more information about RADIUS authentication, see the *ExtremeXOS Concepts Guide*.

Beginning with ExtremeXOS 11.3 you can also specify the destination VLAN to enter upon a successful authentication.

Additional Requirements. This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

User names are not case-sensitive; passwords are case-sensitive. User names must have a minimum of 1 character and a maximum of 32 characters. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you use RADIUS for authentication, Extreme Networks recommends that you use the same user name and password for both local authentication and RADIUS authentication.

If you attempt to create a user name with more than 32 characters, the switch displays the following messages:

```
%% Invalid name detected at '^' marker.
%% Name cannot exceed 32 characters.
```

If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Modifying an Existing Account. To modify an existing local netlogin account, use the following command:

```
configure netlogin local-user <user-name> {vlan-vsa [{tagged | untagged}
[<vlan_name>] | <vlan_tag>]] | none}}
```

Displaying Local Netlogin Accounts. To display a list of local netlogin accounts on the switch, including VLAN information, use the following command:

```
show netlogin local-users
```

Example

The following command creates a local network login user name and password:

```
create netlogin local-user megtest
```

After you enter the local netlogin user name, press [Enter]. The switch prompts you to enter a password (the switch does not display the password):

```
password:
```

After you enter the password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

The following command creates a local network login user name, password, and associates a destination VLAN with this account:

```
create netlogin local-user accounting vlan-vsa blue
```

As previously described, the switch prompts you to enter and confirm the password.

History

This command was first available in ExtremeXOS 11.2.

The `vlan-vsa` parameter and associated options were added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

delete netlogin local-user

```
delete netlogin local-user <user-name>
```

Description

Deletes a specified local netlogin user name and its associated password.

Syntax Description

user-name	Specifies a local netlogin user name.

Default

N/A.

Usage Guidelines

Use the `show netlogin local-users` command to determine which local netlogin user name you want to delete from the system. The `show netlogin local-users` output displays the user name and password in a tabular format.

This command applies only to web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command.

Example

The following command deletes the local network login megtest along with its associated password:

```
delete netlogin local-user megtest
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable netlogin

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Description

Disables Network Login modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of authentication types can be disabled on the same switch. To enable an authentication mode, use the following command:

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables mac-based Network Login:

```
disable netlogin mac
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin dot1x guest-vlan ports

```
disable netlogin dot1x guest-vlan ports [all | <ports>]
```

Description

Disables the guest VLAN on the specified 802.1x netlogin ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
ports	Specifies one or more ports included in the guest VLAN.

Default

Disabled.

Usage Guidelines

Use this command to disable the guest VLAN feature.

Enabling Guest VLANs. To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Example

The following command disables the guest VLAN on all ports:

```
disable netlogin dot1x guest-vlan ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable netlogin logout-privilege

```
disable netlogin logout-privilege
```

Description

Disables Network Login logout window pop-up.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login. When disabled, the logout window pop-up will no longer appear. However, if session refresh is enabled, the login session will be terminated after the session refresh timeout.

Example

The following command disables Network Login logout-privilege:

```
disable netlogin logout-privilege
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin ports

```
disable netlogin ports <ports> [{dot1x} {mac} {web-based}]
```

Description

Disables Network Login on a specified port for a particular method.

Syntax Description

ports	Specifies the ports for which netlogin should be disabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

Netlogin is disabled by default.

Usage Guidelines

Network Login must be disabled on a port before you can delete a VLAN that contains that port.

This command applies to the MAC-based, web-based, and 802.1x mode of network login. To control which authentication mode is used by Network Login, use the following commands:

```
enable netlogin [{dot1x} {mac} {web-based}]  
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables dot1x and web-based Network Login on port 2:9:

```
disable netlogin ports 2:9 dot1x web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin session-refresh

```
disable netlogin session-refresh
```

Description

Disables Network Login session refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Network Login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the LogOut link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to five minutes by default.

This command applies only to the web-based authentication mode of network login.

Example

The following command disables network login session refresh:

```
disable netlogin session-refresh
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Description

Enables Network Login authentication modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies mac-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of types of authentication can be enabled on the same switch. At least one of the authentication types must be specified on the command line.

To disable an authentication mode, use the following command:

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command enables web-based network login:

```
enable netlogin web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin dot1x guest-vlan ports

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Description

Enables the guest VLAN on the specified 802.1x netlogin ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
ports	Specifies one or more ports or slots and ports on which the guest VLAN is enabled.

Default

Disabled.

Usage Guidelines

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Modifying the Supplicant Timer. By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is not a user-configured parameter.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-period <quiet_period>} {reauth-period <reauth_period>} {supp-resp-timeout <supp_resp_timeout>}]
```

Creating the Guest VLAN. Before you can enable the guest VLAN on the specified ports, you must create the guest VLAN. To create the guest VLAN, use the following command:

```
configure netlogin dot1x guest-vlan <vlan_name> {ports <port_list>}
```

Example

The following command enables the guest VLAN on all ports:

```
enable netlogin dot1x guest-vlan ports all
```

The following command enables the guest VLAN on ports 2 and 3:

```
enable netlogin dot1x guest-vlan ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable netlogin logout-privilege

```
enable netlogin logout-privilege
```

Description

Enables network login logout pop-up window.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login logout-privilege:

```
enable netlogin logout-privilege
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin ports

```
enable netlogin ports <ports> [{dot1x} {mac} {web-based}]
```

Description

Enables network login on a specified port for a particular authentication method.

Syntax Description

ports	Specifies the ports for which netlogin should be disabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All methods are disabled on all ports.

Usage Guidelines

For campus mode login with web-based clients, the following conditions must be met:

- A DHCP server must be available, and a DHCP range must be configured for the port or ports in the VLAN on which you want to enable Network Login.
- The switch must be configured as a RADIUS client, and the RADIUS server must be configured to enable the Extreme Network Login capability.

For ISP mode login, no special conditions are required. A RADIUS server must be used for authentication.

Network Login is used on a per port basis. A port that is tagged can belong to more than one VLAN. In this case, Network Login can be enabled on one port for each VLAN.

Windows authentication is not supported via network login.

Example

The following command configures network login on port 2:9 using web-based authentication:

```
enable netlogin ports 2:9 web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin session-refresh

```
enable netlogin session-refresh {<minutes>}
```

Description

Disables network login session refresh.

Syntax Description

minutes	Specifies the session refresh time for network login in minutes.

Default

Disabled, with a value of three minutes for session refresh.

Usage Guidelines

Network Login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default. The value can range from 1 to 255 minutes. When you configure the network login session refresh for the logout window, ensure that the FDB aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

To reset the session refresh value to the default behavior, use this command without the `minutes` parameter.

Example

The following command enables network login session refresh and sets the refresh time to ten minutes:

```
enable netlogin session-refresh 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show banner netlogin

```
show banner netlogin
```

Description

Displays the user-configured banner string for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.

Example

The following command displays the network login banner:

```
show banner netlogin
```

If a custom banner web page exists, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****  
NOTE: Banner is not in use. Overridden since custom login page  
"netlogin_login_page.html" is present.
```

If a custom banner web page does not exist, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show netlogin

```
show netlogin {port <portlist> vlan <vlan_name>} {dot1x {detail}} {mac}
{web-based}
```

Description

Shows status information for Network Login.

Syntax Description

portlist	Specifies one or more ports or slots and ports.
vlan_name	Specifies the name of a VLAN.
dot1x	Specifies 802.1x information.
mac	Specifies MAC-based information.
web-based	Specifies web-based information.
detail	Shows detailed information.

Default

N/A.

Usage Guidelines

Depending on your configuration, software version, and the parameters you choose to display, the information reported by this command may include some or all of the following:

- Whether network login is enabled or disabled.
- The base-URL.
- The default redirect page.
- The logout privileges setting.
- The netlogin session-refresh setting and time.
- The MAC and IP address of supplicants.
- The type of authentication, 802.1x, MAC-based, or HTTP (web-based).
- The guest VLAN configurations, if applicable.
- The dynamic VLAN state and uplink ports, if configured.
- Whether netlogin port restart is enabled or disabled.
- Which order of authentication protocols is currently being used.

If you do not specify the authentication method, the switch displays information for all netlogin authentication methods.

Example

The following command shows the summary network login information:

```
show netlogin
```

The following is sample output from this command:

```
NetLogin Authentication Mode : web-based ENABLED; 802.1x ENABLED; mac-based ENABLED
NetLogin VLAN                : "nvlan"
NetLogin move-fail-action    : Authenticate
NetLogin Client Aging Time   : 5 minutes
Dynamic VLAN Creation        : Enabled
Dynamic VLAN Uplink Ports    : 12
```

----- Web-based Mode Global Configuration -----

```
Base-URL                : network-access.com
Default-Redirect-Page    : http://www.yahoo.com
Logout-privilege         : YES
Netlogin Session-Refresh : ENABLED; 3 minutes
Authentication Database   : Radius, Local-User database
-----
```

----- 802.1x Mode Global Configuration -----

```
Quiet Period                : 60
Supplicant Response Timeout : 30
Re-authentication period    : 200
RADIUS server timeout       : 30
EAPOL MPDU version to transmit : v1
Authentication Database      : Radius
-----
```

----- MAC Mode Global Configuration -----

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:86:3F:1C:35/48	yaqu	any
00:01:20:00:00:00/24	yaqu	any
00:04:0D:28:45:CA/48	=4253C5;500@	any
00:10:14:00:00:00/24	yaqu	any
00:10:A4:A9:11:3B/48	yaqu	any
00:10:A4:00:00:00/24	yaqu	any
Default	yaqu	any

```
Authentication Database      : Radius, Local-User database
-----
```

```
Port: 5, Vlan: nvlan, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled
```


MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 9, Vlan: nvlan, State: Enabled, Authentication: web-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 10, Vlan: nvlan, State: Enabled, Authentication: 802.1x, mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 17, Vlan: engr, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 17, Vlan: mktg, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 19, Vlan: corp, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No		0	00040D50E13A
00:10:dc:98:54:00	10.201.31.113	Yes, Radius	802.1x	24	md5isp7

Port: 19, Vlan: nvlan, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No	802.1x	0	

Port: 19, Vlan: voice-ip, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	Yes, Radius	802.1x	75	00040D50E13A

The following command shows more detailed information, including the configured authentication methods:

```
show netlogin port 3:2 vlan "Default"
```

The following is sample output from this command:

```
Port: 2:1      Vlan: Default
Authentication: Web-Based, 802.1x
Port State:    Unauthenticated
Guest VLAN:    Not Enabled
DHCP:          Not Enabled
```

MAC	IP address	Auth	Type	ReAuth-Timer	User
00:0C:F1:E8:4E:13	0.0.0.0	No	802.1x	0	Unknown
00:01:30:F3:EA:A0	10.0.0.1	Yes	802.1x	0	testUser

The following command shows information about a specific port configured for netlogin:

```
show netlogin port 1:1
```

The following is sample output from this command:

```
Port      : 1:1
Port Restart : Enabled
Vlan      : Default
Authentication: mac-based
Port State : Enabled
Guest Vlan : Disabled
```

MAC	IP address	Auth	Type	ReAuth-Timer	User

The following command shows the details of the 802.1x mode:

```
show netlogin dot1x detail
```

The following is sample output from this command:

```
NetLogin Authentication Mode : web-based DISABLED; 802.1x ENABLED; mac-based
DISABLED
NetLogin VLAN                  : "nl"
NetLogin move-fail-action      : Deny
```

```
-----
                        802.1x Mode Global Configuration
-----
Quiet Period                : 30
Supplicant Response Timeout : 30
Re-authentication period    : 3600
RADIUS server timeout       : 30
EAPOL MPDU version to transmit : v1
Guest VLAN                  : destVlan
-----
```

```
Port: 1:1, Vlan: Default, State: Enabled, Authentication: 802.1x, Guest Vlan:
destVlan
```

```

MAC
00:00:86:53:c3:14 : IP=0.0.0.0      Auth=Yes User= testUser
                   : AuthPAE state=AUTHENTICATED BackAuth state=IDLE
                   : ReAuth time left=3595      ReAuth count=1
                   : Quiet time left=37
00:02:03:04:04:05 : IP=0.0.0.0      Auth=No  User=
                   : AuthPAE state=CONNECTING  BackAuth state=IDLE
                   : ReAuth time left=0        ReAuth count=2
                   : Quiet time left=37
-----

```

For 802.1x, if re-authentication is disabled, the re-authentication period appears as follows:

```
Re-authentication period : 0 (Re-authentication disabled)
```

The following command:

```
show netlogin port 5:4 dot1x
```

Generates the following sample output:

```

Port          : 5:4
Port Restart  : Disabled
Vlan          : corp
Authentication : 802.1x
Port State    : Enabled
Guest Vlan    : Enabled

```

```

MAC          IP address  Authenticated  Type    ReAuth-Timer  User
00:10:dc:92:53:2d  10.201.31.119  Yes,Radius    802.1x   14            md5isp4
-----

```

The following command:

```
sh netlogin port 5:4 dot1x detail
```

Generates the following sample output:

```

Port          : 5:4
Port Restart  : Disabled
Vlan          : corp
Authentication : 802.1x
Port State    : Enabled
Guest Vlan    : Enabled

```

```

MAC
00:10:dc:92:53:2d : IP=10.201.31.119 Auth=Yes User=md5isp4
                   : AuthPAE state=AUTHENTICATED BackAuth state=IDLE
                   : ReAuth time left=8      ReAuth count=0
                   : Quiet time left=0
-----

```

History

This command was first available in ExtremeXOS 11.1.

Information about the guest VLAN was added in ExtremeXOS 11.2.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Information about dynamic VLANs and netlogin port restart was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show netlogin guest-vlan

```
show netlogin guest-vlan {vlan_name}
```

Description

Displays the configuration for the guest VLAN feature.

Syntax Description

vlan_name	Specifies the name of a guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the guest VLANs configured on the switch.

If you specify the `vlan_name`, the switch displays information for only that guest VLAN.

The output displays the following information in a tabular format:

- `Port`—Specifies the 802.1x enabled port configured for the guest VLAN.
- `Guest-vlan`—Displays the enabled/disabled state of the guest VLAN feature.
- `Vlan`—Specifies the name of the guest VLAN.

Example

The following command displays the local netlogin list:

```
show netlogin guest-vlan
```

The following is sample output from this command:

Port	Guest-vlan	Vlan
5:1	Disabled	gv11
5:2	Enabled	gv12
5:3	Disabled	gv13
5:4	Enabled	gv14

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show netlogin local-users

```
show netlogin local-users
```

Description

Displays the local netlogin users configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the list of local netlogin users and associated VLANs.

If you associated a VLAN with a local netlogin user, the output displays the name of the VLAN. If you have not associated a VLAN with a local netlogin user, the output displays not configured.

The Extended-VLAN VSA column displays the name of the VLAN and the following information:

- <not configured>—Specifies that you have not associated a VLAN with a local netlogin user.
- *—Specifies the movement based on the incoming port's traffic. For example, the VLAN behaves like VSA 203 if identified with a VLAN name or VSA 209 if identified with a VLAN ID.
- T—Specifies a tagged client.
- U—Specifies an untagged client.

In addition, this output is useful to determine which local netlogin user you want to modify or delete from the system.

Example

The following command displays the local netlogin list:

```
show netlogin local-users
```

The following is sample output from this command:

Netlogin Local User Name	Password (encrypted)	Extended-VLAN VSA
000000000012	Iqyydz\$MP7AG.VAmwOoqiKX2u13H1	U hallo
00008653C314	BoO28L\$oRVvKv8.wmxcorhhXxQY40	* default
megtest	w7iMbp\$lBL34/dLx4G4M8aAdiCvI	<not configured>
testUser	/Jhouw\$iHE15steebwhOibgj6pZq.	T testVlan

History

This command was first available in ExtremeXOS 11.2.

The output was modified to include VLAN information in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show netlogin mac-list

```
show netlogin mac-list
```

Description

Displays the MAC address list for MAC-based network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the MAC address list used for MAC-based network login.

MAC-based authentication is VR aware, so there is one MAC list per VR.

Example

The following command displays the MAC address list:

```
show netlogin mac-list
```

The following is sample output from this command:

MAC Address/Mask	Password (encrypted)	Port(s)
-----	-----	-----
00:00:00:00:00:10/48	<not configured>	1:1-1:5
00:00:00:00:00:11/48	<not configured>	1:6-1:10
00:00:00:00:00:12/48	<not configured>	any
00:01:30:70:0C:00/48	yaqu	any
00:01:30:32:7D:00/48	ravdqsr	any
00:04:96:00:00:00/24	<not configured>	any

History

This command was first available in ExtremeXOS 11.1.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

unconfigure netlogin authentication database-order

```
unconfigure netlogin [mac | web-based] authentication database-order
```

Description

Restores the default order of database authentication protocols to use.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

By default, the authentication order is Radius, Local-User database.

Usage Guidelines

Use this command to restore the default configuration order for the database authentication protocols. For details see [configure netlogin authentication database-order](#).

Example

The following command sets the database authentication order to Radius, Local user database for mac-based authentication:

```
unconfigure netlogin mac authentication database-order
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

unconfigure netlogin dot1x guest-vlan

```
unconfigure netlogin dot1x guest-vlan {ports <port_list> | <vlan_name>}
```

Description

Unconfigures the guest VLAN feature for 802.1x authentication.

Syntax Description

ports_list	Specifies one or more ports included in the guest VLAN.
vlan_name	Specifies all ports included in the guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the guest VLAN for 802.1x authentication.

If you do not specify one or more ports or the VLAN name, this command unconfigures all of the 802.1x ports configured for the guest VLAN feature.

If you specify one or more ports, this command unconfigures the specified 802.1x ports for the guest VLAN feature.

If you specify the VLAN name, this command unconfigures all of the 802.1x ports configured for the specified guest VLAN.

Example

The following command unconfigures the guest VLAN feature for 802.1x authentication:

```
unconfigure netlogin dot1x guest-vlan
```

History

This command was first available in ExtremeXOS 11.2.

The `ports` option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

unconfigure netlogin vlan

```
unconfigure netlogin vlan
```

Description

Unconfigures the VLAN for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled.

Example

The following command unconfigures the network login VLAN:

```
unconfigure netlogin vlan
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

This chapter describes commands for:

- Enabling and disabling CLEAR-Flow
- Displaying CLEAR-Flow rules
- Displaying triggered CLEAR-Flow rules

CLEAR-Flow is a broad framework for implementing security, monitoring, and anomaly detection in ExtremeXOS software. Instead of simply looking at the source and destination of traffic, CLEAR-Flow allows you to specify certain types of traffic that require more attention. Once certain criteria for this traffic are met, the switch can either take an immediate, pre-determined action, or send a copy of the traffic off-switch for analysis.

CLEAR-Flow is an extension to Access Control Lists (ACLs). You create ACL policy rules to count packets of interest. CLEAR-Flow rules are added to the policy to monitor these ACL counter statistics. The CLEAR-Flow agent monitors the counters for the situations of interest to you and your network. You can monitor the cumulative value of a counter, the change to a counter over a sampling interval, the ratio of two counters, or even the ratio of the changes of two counters over an interval. For example, you can monitor the ratio between TCP SYN and TCP packets. An abnormally large ratio may indicate a SYN attack.

If the rule conditions are met, the CLEAR-Flow actions configured in the rule are executed. The switch can respond by installing an ACL that will block or rate limit the traffic, executing a set of CLI commands, or sending a report using a SNMP trap or EMS log message.

**NOTE**

CLEAR-Flow is available only on the BlackDiamond 10808 family and BlackDiamond 12800 series switches.

disable clear-flow

```
disable clear-flow
```

Description

Disable the CLEAR-Flow agent.

Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is disabled, sampling stops and the and all rules are left in the current state. It will not reset actions that were taken while CLEAR-Flow was enabled.

Example

The following example disables CLEAR-Flow on the switch:

```
disable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

enable clear-flow

```
enable clear-flow
```

Description

Enable the CLEAR-Flow agent.

Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is enabled, sampling begins and actions are taken based on the CLEAR-Flow rules that are configured on the switch.

Example

The following example enables CLEAR-Flow on the switch:

```
enable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

show clear-flow

```
show clear-flow
```

Description

Displays the status of the CLEAR-Flow agent, any CLEAR-Flow policies on each interface, and the number of CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following display shows output for the command `show clear-flow`:

```
clear-flow: Enabled
```

VLAN	Port	Policy Name	No. of CF Rules
*	2:1	CFexample	6
*	2:26	CFexample	6
*	2:40	CFexample	6
Default	*	CFexample	6

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

show clear-flow acl-modified

```
show clear-flow acl-modified
```

Description

Displays the ACLs modified by CLEAR-Flow actions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ACLs that have been modified by CLEAR-Flow rules that have been triggered.

Example

The following display shows output for the command `show clear-flow acl-modified`:

Policy Name	Vlan Name	Port	Rule Name	Default ACL Actions	CF Added Actions
clearFlow	*	2:26	acl-rule-4	D	QP1
clearFlow	*	2:26	acl-rule-3	D	D
clearFlow	*	2:26	acl-rule-2	D	M
clearFlow	*	2:26	acl-rule-1		P
clearFlow	Default	*	acl-rule-4	D	QP1
clearFlow	Default	*	acl-rule-3	D	D
clearFlow	Default	*	acl-rule-2	D	M
clearFlow	Default	*	acl-rule-1		P

=====
Total Entries: 8
Notation:
P - Permit, D- Deny, M - mirror enabled, m - mirror disabled

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

show clear-flow rule

```
show clear-flow [port <port> | vlan <vlanname> | any] {rule <rulename>}
{detail}
```

Description

Displays the CLEAR-Flow rules, values, and configuration.

Syntax Description

port	Specifies the port.
vlanname	Specifies the VLAN.
any	Specifies the wildcard interface.
rulename	Specifies the entry name of a CLEAR-Flow rule.
detail	Display detailed information.

Default

N/A.

Usage Guidelines

If you issue the command without the `rule` keyword, all of the CLEAR-Flow rules for the policy on the port, VLAN, and the wildcard are displayed. If you specify a rule name, only that rule will be displayed. The `detail` keyword displays detailed information about the rule.

Example

The following display shows output for the command `show clear-flow port 2:6`:

Rule Name	Type	Period	Last Value	Rel Oper	Threshold	TCNT	NumAction	If	Else
rule-count	CN	30	16892762	>	100	7	3	3	
rule-delta	DT	30	7762385	>	1000	1	4	3	
rule-delta-2	DT	5	0	>	1000	0	4	3	
rule-delta-ratio	DR	30	0	>	20	0	2	0	
rule-ratio	RT	30	0	>	10	0	3	3	
rule-ratio-2	RT	5	0	>	10	0	3	3	

Total Entries: 6

Notation:

Threshold Type: CN - Count, DT - Delta, RT - Ratio, DR - DeltaRatio

TCNT - Number of times expression is continuously evaluated to be true

The following display shows output for the command `show clear-flow port 2:6 rule rule-delta detail`:

```
Rule Name: rule-delta      Sample Period: 30      Hysteresis: 20
=====
    DELTA(counter1) = 0 sampled at 24 seconds ago
    Expression evaluation is currently FALSE

if (DELTA(counter1) > 1000) then {
    PERMIT:    Allow ACL rule acl-rule-3
    SYSLOG:    [INFO] [Delta $ruleValue counter $counter1 offset $counterOffset1
delTime $deltaTime delay $delayTime]
    CLI:       [disable port $port]
    QOS:       Set rule acl-rule-4 qos value to QP6
} else {
    DENY:     Block ACL rule acl-rule-3
    QOS:       Set rule acl-rule-4 qos value to QP1
    CLI:       [enable port $port]
}
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

show clear-flow rule-all

```
show clear-flow rule-all
```

Description

Displays all the CLEAR-Flow rules on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following display shows output for the command `show clear-flow rule-all`:

Policy Name	Vlan Name	Port	Rule Name	Last Value	OP	Threshold	TCNT	Sec
clearFlow	*	2:1	rule-count	1	>	100	0	11
clearFlow	*	2:1	rule-delta	1	>	1000	0	11
clearFlow	*	2:1	rule-delta	0	>	1000	0	4
clearFlow	*	2:1	rule-delta	0	>	20	0	11
clearFlow	*	2:1	rule-ratio	0	>	10	0	11
clearFlow	*	2:1	rule-ratio	0	>	10	0	4
clearFlow	*	2:26	rule-count	9030635	>	100	1	10
clearFlow	*	2:26	rule-delta	9030635	>	1000	1	10
clearFlow	*	2:26	rule-delta	0	>	1000	0	4
clearFlow	*	2:26	rule-delta	0	>	20	0	10
clearFlow	*	2:26	rule-ratio	0	>	10	0	10
clearFlow	*	2:26	rule-ratio	0	>	10	0	4
clearFlow	Default	*	rule-count	36666439	>	100	1	10
clearFlow	Default	*	rule-delta	36666439	>	1000	1	10
clearFlow	Default	*	rule-delta	0	>	1000	0	4
clearFlow	Default	*	rule-delta	0	>	20	0	10
clearFlow	Default	*	rule-ratio	0	>	10	0	10
clearFlow	Default	*	rule-ratio	0	>	10	0	4

=====
Total Entries: 18
Notation:
TCNT - Number of times expression is continuously evaluated to be true
Sec - Number of seconds elapsed from last sampled data

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

show clear-flow rule-triggered

```
show clear-flow rule-triggered
```

Description

Displays the triggered CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the rules that have been triggered; in other words, the rule threshold has been reached.

Example

The following display shows output for the command `show clear-flow rule-triggered`:

Policy Name	Vlan Name	Port	Rule Name	Last Value	OP	Threshold	TCNT	Sec
clearFlow	*	2:26	rule-count	9130377	>	100	2	25
clearFlow	*	2:26	rule-delta	99742	>	1000	2	25
clearFlow	Default	*	rule-count	37069465	>	100	2	25
clearFlow	Default	*	rule-delta	403026	>	1000	2	25

Total Entries: 4

Notation:

TCNT - Number of times expression is continuously evaluated to be true

Sec - Number of seconds elapsed from last sampled data

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 10808 and BlackDiamond 12800 series platforms only.

This chapter describes commands for:

- Configuring SummitStack
- Enabling and disabling SummitStack
- Deploying SummitStack
- Displaying SummitStack Details

SummitStack

SummitStack allows you to physically connect up to eight individual Summit switches together as a single logical unit. This logical unit behaves as a single switch with a single IP address and a single point of authentication.

In ExtremeXOS, a stack is controlled by a master switch. The master runs full ExtremeXOS and is responsible for maintaining all of the software tables for all the switches in the stack. There can only be one master in a stack of switches. All switches in the stack are referred to as stack members, stackables, switches, or nodes.

An easy way to understand a stack is to think of a stack as a virtual chassis. Each node acts as if it were a slot in a chassis and waits for configuration information from the master. The high-speed stacking links behave like the backplane links of a chassis.

SummitStack is supported on the following devices:

- X450e-24p
- X450a-24t
- X450a-24tDC
- X450a-48t
- X450a-24x
- X450a-24xDC
- X450e-48p
- X450a-48tDC
- X450-24t
- X450-24x
- X250e-24x
- X250e-24p
- X250e-24t
- X250e-48p
- X250e-48t

configure stacking alternate-ip-address

```
configure stacking alternate-ip-address [<ipaddress> <netmask> |
<ipNetmask>] <gateway> automatic

configure stacking [node-address <node-address> | slot <slot-number>]
alternate-ip-address [<ipaddress> <netmask> | <ipNetmask>] <gateway>
```

Description

Configures an alternate management IP address and subnetwork

Syntax Description

node-address	Specifies the MAC address assigned to the target stackable by the factory. A node address or slot number is required unless the automatic keyword is specified.
slot-number	Specifies a number between 1 and 8 that selects a node in the active topology.
ipaddress netmask	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node together with the subnetwork mask specified for the Management subnetwork. Example: 66.77.88.1 255.255.255.0.
ipNetmask	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node, followed by a slash (/) character, followed by a decimal number that represents the number of leading one bits in the subnetwork address. An example is 66.77.88.1/24.
gateway	The address of an IP router. A default route is set up to reach this gateway.

Default

No alternate IP address is configured.

Usage Guidelines

If a Mgmt subnetwork is configured and the alternate IP subnetwork does not exactly match the configured Mgmt subnetwork, the information configured by one of the commands specified above is not used. The previously configured alternate IP address is removed if it was installed and subsequently a Mgmt subnetwork is configured that does not exactly match the alternate IP subnetwork. In either case, an error message is logged. The alternate IP address is used if there is no configured Mgmt subnetwork.

To use the command with the node address, the node must be in the stack topology. This command operates only on one node at a time. There are no checks to verify that the address is the one configured in the management VLAN subnet.

The command that does not require a node address specifies the automatic keyword. Usage of this command causes an alternate IP address to be assigned to every node in the stack topology. The first address is the address specified in the [<ipaddress> <netmask> | <ipNetmask>] parameter. The next address is the IP address plus one, and so on. Since there is a specified subnet mask, the address will be checked to insure that the block of IP addresses fits within the specified subnet given the number of nodes in the stack topology. The address block is tested to insure that each address is a valid IP unicast

address. If the test fails, no node is configured and an error message is printed. Assignment is in the order in which nodes would currently appear in the [show stacking](#) display.

The configuration takes effect after the command is successfully executed.

The alternate IP address, subnetwork, and gateway are not used when the node is not operating in stacking mode.

Example

To configure an alternate IP address for every node in the stack with a single command:

```
configure stacking alternate-ip-address 10.120.1.10/24 10.120.1.1 automatic
```

To configure an alternate IP address on a single node in the stack topology:

```
configure stacking node-address 00:04:96:26:6b:ed alternate-ip-address 10.120.1.1/24  
<gateway>
```

You may configure an alternate IP address using a slot number for a node that is currently occupying the related slot:

```
configure stacking slot 4 alternate-ip-address 10.120.1.13/24 10.120.1.1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking easy-setup

```
configure stacking easy-setup
```

Description

This command provides an easy way to initially configure the stacking parameters of all nodes in a new stack.

Syntax Description

This command does not have additional syntax.

Default

N/A

Usage Guidelines

This command performs the following functions:

- Informs you of the stacking parameters that will be set.
- Informs you of the number of nodes that will be configured.
- Inform you whether minimal or no redundancy will be configured, and which slot will contain the master node.
- Inform you of the slot number that will be assigned to the node on which your management session is being run.
- If applicable, warns you that the current configuration file changes will be lost and you need to save the files.
- If a daisy chain is present, warns you that you should wire the stack as a ring before running this command.
- Requires you to confirm before the operation takes place. If you proceed, the command will
 - Enable stacking on all nodes.
 - Configures the stacking MAC address using the factory address of the current node.
 - Configure a slot number for each node.
 - Configure redundancy to minimal in a ring topology or none in a daisy chain topology.
 - Reboot the stack topology.

Enable stacking is performed as if the “enable stacking” command was issued.

Configuring the stacking mac-address is performed as if the `configure stacking mac-address` was issued on the current node.

Configuring the stacking slot numbers is performed as if the `configure stacking slot-number automatic` command was issued on the current node.

On a daisy chain topology, configuring the master-capability is performed as if the `configure stacking redundancy none` command was issued. On a ring topology, configuring the master-capability is performed as if the `configure stacking redundancy minimal` command was issued.

If you choose not to proceed with the setup, the following message is displayed:

```
Cancelled easy stack setup configuration.
```

Example

If you have an 8-node stack in a ring topology and has powered on all the nodes, the [show stacking](#) command will show the stack topology as ring and that all intended nodes are present. If you have not changed any ExtremeXOS configuration, the command displays:

```
* X450a-24t.30 # configure stacking easy-setup
```

For every node in the 8-node stack, this command will:

- enable stacking
- configure a stack MAC address
- choose and configure a slot number (this node will be assigned to slot 1)
- configure redundancy to minimal (slot 1 will be the Master node)

Upon completion, the stack will automatically be rebooted into the new configuration.

Warning: If stacking is already configured, this command will alter that configuration.

Warning: There are unsaved configuration changes. You may wish to save them before proceeding.

Do you wish to proceed? (y/N) y

Stacking configuration is complete. Rebooting...

If the 8-node stack topology is a daisy chain, and the user is logged into a node in the middle of the chain, the command output might appear as follows:

```
* X450a-24t.30 # configure stacking easy-setup
```

For every node in the 8-node stack, this command will:

- enable stacking
- configure a stack MAC address
- choose and configure a slot number (this node will be assigned to slot 5)
- configure redundancy to none (slot 1 will be the master node)

Upon completion, the stack will automatically be rebooted into the new configuration.

Warning: If stacking is already configured, this command will alter that configuration.

Warning: This stack is a daisy chain. It is highly recommended that the stack be connected as a ring before running this command.

Do you wish to proceed? (y/N) Yes

Stacking configuration is complete. Rebooting...

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking license-level

```
configure stacking {node-address <node-address> | slot <slot-number>}
license-level [core | advanced-edge | edge]
```

Description

Allows you to restrict the license level at which the node operates.

Syntax Description

node-address	Specifies the MAC address assigned to a stackable by the factory.
slot-number	Specifies the number of the slot currently occupied by the node on which you need to configure the license level.

Default

No license level restriction is configured.

Usage Guidelines

This command restricts a node to operating at a lower license level than the level that was purchased for the node.

Running this command will not change the installed license level. For example, if a stackable is configured with the Advanced Edge license and you configure a license level restriction of Edge, the unit will be restricted to features available in the Edge license. However, you will be able to remove the restriction and operate at the Advanced Edge level.

If the installed license level of the target node is lower than the level you are attempting to configure, the following message appears:

Warning: Switch will not operate at a license level beyond that which was purchased.

If the node-address or slot parameter is not specified, the command takes effect on every node in the stack topology.

This command takes effect after you restart the node. The following message appears after the command is executed:

This command will take effect at the next reboot of the specified node(s).

If you restart the node without configuring a license level restriction, the node operates at the purchased license level. To see the purchased license level of a node, run show license after logging in to the node.

The `show license` command displays the current license level in use as the “Effective License Level”:

```
Slot-2 Stack.1 # show licenses
Enabled License Level:
    Advanced Edge
Enabled Feature Packs:
    None
Effective License Level:
    Edge
```

The `show stacking configuration` and `show stacking detail` commands allow you to see the configured license level restriction and the restriction currently in use.

The effective license level will appear only when stacking is enabled. The command is node-specific. The effective license level is the level at which the node is restricted to operate, and is not necessarily the level at which the entire stack is operating. This is because it is possible to have the restriction differ on each node, in which case one or more nodes may have failed because of the differing levels.

Example

To configure stacking level Edge on all nodes on a stack:

```
configure stacking license-level edge
```

To configure stacking level Edge for a node:

```
configure stacking node-address 00:04:96:26:6b:ed license-level edge
```

To configure stacking level Advanced Edge for an active node that currently occupies slot 4:

```
configure stacking slot 4 license-level advanced-edge
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking mac-address

```
configure stacking {node-address <node-address> | slot <slot-number>} mac-address
```

Description

Selects a node in the stack whose factory assigned MAC address is to be used to form the stack MAC address. The formed address is then configured on every node in the stack topology.

Syntax Description

·node-address	Specifies a node that is present in the stack topology using its factory assigned MAC address.
·slot-number	Specifies a node that is present in the active topology using its slot number.

Default

No stacking MAC selection is configured.

Usage Guidelines

You must select a node whose factory assigned MAC address will be used to form a MAC address that will represent the stack as a whole. The system forms the stack MAC address by setting the Universal / Local bit in the specified MAC address. This means that the stack MAC address is a locally administered address, and not the universal MAC address assigned to the selected node.

If you do not specify any node, the selected node is the node from which you are running the command. All nodes in the stack topology are configured with the formed stack MAC address.

This command takes effect only after you restart the node. The following message appears after you run the command:

```
This command will take effect at the next reboot of the specified node(s).
```

If a stack node that has just joined the stack detects that its stack MAC address is not configured or is different than the stack MAC address in use, it will log the following message at the Error log level:

```
The stack MAC address is not correctly configured on this node. The stack can not operate properly in this condition. Please correct and reboot.
```

If you have not configured (or inconsistently configured) the stack MAC address you might encounter difficulty in diagnosing the resulting problems. Whenever the master node (including itself) detects that one or more nodes in its active topology do not have the correct or any stack MAC address configured, it will display the following message to the console every five minutes until you configure a MAC address and restart the node(s):

```
The stack MAC address is either not configured or its configuration is not consistent within the stack. The stack can not operate properly in this condition. Please correct and reboot.
```

Example

To select the node to which you have logged in to supply the MAC address for stack MAC address formation:

```
configure stacking mac-address
```

To select a node other than the one to which you are logged in to supply the MAC address for stack MAC address formation:

```
configure stacking node-address 00:04:96:26:6b:ed mac-address
```

To select an active node to supply the MAC address for stack MAC address formation:

```
configure stacking slot 4 mac-address
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking master-capability

```
configure stacking [node-address <node-address> | slot <slot-number>]
master-capability [on | off]
```

Description

The command configures a node to be allowed to operate as either a backup or master, or prevents a node from operating as either.

The command controls the setting on the specified node only. To set the master capability for all nodes on a stack, you can use the command `configure stacking redundancy`.

Syntax Description

node-address	Specifies the MAC address assigned to the stackable by the factory.
slot-number	Specifies a number between 1 and 8 that selects a node in the active topology as shown when the <code>show stacking</code> command is executed.

Default

Master-capability is On.

Usage Guidelines

At least one node in the stack topology must be master-capable.

If you attempt to disable the master-capability of the only master capable node in a stack topology, the attempt is denied and following message appears:

```
Error: At least one node must have Master-capability configured "on".
```

This command is used to set up master-capability manually. It can also be used to adjust the result achieved when the `configure stacking redundancy` command is used.

The setting takes effect the next time the node reboots. When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To turn on the master capability for a node:

```
configure stacking node-address 00:04:96:26:6b:ed master-capability on
```

To turn on the master capability of an active node currently occupying slot 4:

```
configure stacking slot 4 master-capability on
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking priority

```
configure stacking {node-address <node-address> | slot <slot-number>}
priority [<node-pri> | automatic]
```

Description

Configures a priority value to be used to influence master and backup election.

Syntax Description

node-address	Specifies the MAC address assigned to the stackable by the factory.
slot-number	Specifies a number between 1 and 8 that selects a node in the active topology as shown when the show stacking command is executed.
node-pri	Specifies the priority as a value between 1 and 100.

Default

No priority is configured.

Usage Guidelines

The node role election priority is a value that is internally calculated by ExtremeXOS for each node. This calculated value determines the nodes that are elected as master and backup. This command allows you to configure a priority value that affects the outcome of this calculation. You can configure the priority on any node in a stack topology. You can specify an integer node-pri value between 1 and 100. The larger the value, the greater the node role election priority.

If no node address or slot is specified, the command takes effect on all nodes at the next node role election cycle. Priority configuration has no operational effect on switches that are not in stacking mode.

If configured on every node, automatic priority commands ExtremeXOS to determine the node role election priority of each active node. Currently, the automatic priority algorithm chooses the master-capable node with the lowest slot number as master and the node with the second lowest slot number as backup. Extreme networks may alter this behavior in later releases.

If you have configured a node with automatic priority and if you have configured another node to use a node-pri value, the node with automatic priority will use zero as the node-priority value during the node role election.

Example

To allow ExtremeXOS to determine node role election priority:

```
configure stacking priority automatic
```

To configure the node priority for the stackable in slot 4:

```
configure stacking slot 4 priority 50
```

To configure the automatic priority algorithm for the stackable with node address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed priority automatic
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking redundancy

```
configure stacking redundancy [none | minimal | maximal]
```

Description

This command sets a master-capability value for each node in the stack topology.

Syntax Description

none	Only one node will have master-capability turned on and all other nodes will have master-capability turned off.
minimal	Two nodes will have master-capability turned on and all other nodes will have master-capability turned off.
maximal	All nodes will have master-capability turned on.

Default

Default value is maximal.

Usage Guidelines

To run this command, the stack should not contain more than eight nodes in its stack topology.

If there are more than eight nodes in the stack topology, the following message appears and the command is not executed:

```
ERROR: This command can only be used when the stack has eight nodes or less.
```

Since only eight nodes can be operational in an active topology at a time, you must disconnect the remaining nodes before configuring master-capability using this command.

If you are using the none or minimal redundancy configuration:

- The configured values of slot-number and priority decide the nodes on which the master-capability should be turned on.
- If the priority values are configured on the nodes, the highest priority node(s) will be chosen.
- If the priority values of all nodes are set to automatic or to the same priority value, the node(s) with the lowest slot number(s) will be chosen. Extreme Networks may change automatic priority behavior in a future release.

If there is a slot number tie or if neither slot-number nor priority were ever configured, the following message appears and the command is not executed:

```
ERROR: Unique slot numbers must be configured before using this command.
```

The setting takes effect at the time of the next restart of the node. The following message appears after the command is successfully executed:

```
This command will take effect at the next reboot of the specified node(s).
```

Redundancy configuration has no operational effect on a node that is not in stacking mode.

Example

To turn on master-capability on all nodes:

```
configure stacking redundancy maximal
```

To turn on master-capability on only one node:

```
configure stacking redundancy none
```

To turn on master-capability on two nodes:

```
configure stacking redundancy minimal
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

configure stacking slot-number

```
configure stacking slot-number automatic
```

```
configure stacking node-address <node-address> slot-number <slot-number>
```

Description

Configures a slot number on one or all nodes in the stack topology.

Syntax Description

node-address	specifies the factory assigned MAC address of the node whose slot number is to be configured.
slot-number	specifies a number between 1 and 8 that is to be assigned as the slot number of the stackable to be used when the stackable is in the active topology.

Default

The default slot-number for a node in stacking mode is 1.

Usage Guidelines

The configuration is stored on the affected node(s) immediately but does not take effect until the next reboot of the node(s). The configuration has no operational effect when the node is not running in stacking mode. To see the configured and active slot numbers of all nodes, use the `show stacking configuration` command.

If a node-address and a slot number are specified, then the node is configured with the specified slot number. There is no check for duplicate slot number at this time; the number is simply assigned as requested.

If the automatic keyword is specified, then automatic slot number assignment is selected, and the assignment is performed on all nodes in the stack topology. If there are more than eight nodes in the stack topology, the assignment is only performed to the first eight nodes.

Automatic slot number assignment causes assignment of slot numbers starting from 1 and increasing up to 8. The nodes in the stack topology are assigned the numbers in the order in which they would appear currently in the `show stacking` command output. In a ring, slot number 1 is assigned to the current node, slot number 2 is assigned to the node connected to the current node's stack port 2, and so forth. In a daisy chain, slot 1 is assigned to the node at the end of the chain that begins with the node connected to the current node's stack port 1.

To see the resulting slot number assignment, run the `show stacking configuration` command.



NOTE

Failure to configure a node will not prevent configuration of the slot numbers on the other nodes, and will not affect the slot number assigned to each node.

Before the command is executed when the automatic option was specified, the following confirmation is solicited:

```
Reassignment of slot numbers may make the stack incompatible with the current
configuration file.
Do you wish to continue? (y/n)
```

When this command is executed successfully, the following message is displayed:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To configure all slot-numbers for a stack:

```
configure stacking slot-number automatic
```

To configure the slot number 4 for the node with MAC address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed slot-number 4
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

disable stacking

```
disable stacking {node-address <node-address>}
```

Description

This command disables the stacking on one or all nodes in the stack topology.

Syntax Description

node address	The MAC address of the node on which you need to disable stacking.

Default

Default value is stacking disabled.

Usage Guidelines

If you do not specify the node-address, stacking is disabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is disabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

A node in the stack topology that is disabled for stacking will not forward the customer's data through its stacking links and will not become a member of the active topology.

A disabled node will become its own master and will process and execute its own configuration independently.

When this command is executed successfully, the following message is displayed:

```
This command will take effect at the next reboot of the specified node(s).
```

Use `show stacking configuration` command to see the current configuration of the stack. Verify the flags in `show stacking configuration` output to confirm that the stacking is disabled on the specified node(s).

Example

The following example disables stacking on an 8 node stack:

```
* X450a-24x.3 # disable stacking
```

```
This command will take effect at the next reboot of the specified node(s).
```

The following example disables stacking on node 5, with a MAC address 00:04:96:26:6b:ed:

```
* X450a-24x.3 # disable stacking node-address 00:04:96:26:6b:ed
```

```
This command will take effect at the next reboot of the specified node(s).
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on Summit platforms that supports SummitStack.

enable stacking

```
enable stacking {node-address <node-address>}
```

Description

This command enables stacking on one or all nodes.

Syntax Description

node-address	The MAC address of the node on which you want to enable stacking.

Default

Default value is stacking disabled.

Usage Guidelines

This command enables stacking on one or all nodes. When a node is operating in stacking mode, QoS profiles QP6 and QP7 cannot be created.

If a node-address is not specified, this command first performs an analysis of the current stacking configuration on the entire stack. If the stack has not yet been configured for stacking operation, or if the configuration is self-inconsistent, the user is offered the option of invoking the easy setup function. The following message appears:

```
You have not yet configured all required stacking parameters.
Would you like to perform an easy setup for stacking operation? (y/N)
```

If you enter Yes, the easy setup procedure is invoked and you first see the following message:

```
Executing "configure stacking easy-setup" command...
```

If you enter No, you will see the following message.

```
Stacking has been enabled as requested.
```

The following describes the operation performed if easy setup is neither offered nor selected.

If you do not enter any node-address, stacking is enabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is enabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

The `show stacking configuration` command shows the current configuration of this parameter as well as the value currently in use.

A node that is enabled for stacking will attempt to join the active topology. If successful, it will then negotiate a node role with the other nodes in the stack and will then become an operational node in the stack according to its role. The master node's configuration will be applied to the node.

When this command is executed successfully, the following message is displayed:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To enable stacking on a stack:

```
* X450a-24x.3 # enable stacking
```

This command will take effect at the next reboot of the specified node(s).

To enable stacking on node 5, with a MAC address 00:04:96:26:6b:ed:

```
* X450a-24x.3 # enable stacking node-address 00:04:96:26:6b:ed
```

This command will take effect at the next reboot of the specified node(s).

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

show power (SummitStack Only)

```
show power
```

Description

Displays the number of power modules present and providing power in each slot.

Syntax Description

This command has no arguments or variables.

Default

Default value

Usage Guidelines

This command is available on all platforms. However, it produces completely different output on a SummitStack. The following table is applicable only when the node is active in the related slot.

Power Supply	Flag	Meaning
Internal	F	Failed or no Power
Internal	P	Power available
Internal	O	48V powered off (48p Summits only)
External (non 48P)	-	Empty
External (non 48P)	F	Failed or no power
External (non 48P)	P	Power available
External (non 48P)		Power supply can never occupy this position
External (48P only)	-	Empty or no power to all PSUs present
External (48P only)	F	Failed or no power (at least 1 PSU has power)
External (48P only)	P	Power available

The Summit X450e-48p and X250e-48p switches accept an external power chassis that holds up to three power supplies. All other Summit family switches accept an external power chassis that holds only one power supply. For Summit family switches other than the Summit X450e-48p and X250e-48p switches, the External PSU columns are left blank.

If a node is not present in a slot, the slot number is printed and the remainder of the row is blank.

On the X450e-48p and X250e-48p external power supply unit, ExtremeXOS cannot identify specific power slots that have power supplies installed. If one slot is empty, the right-most External PSU column will show as Empty (-). If two slots are empty, right-most two External PSU columns will show as Empty (-).

On the X450e-48p and X250e-48p, software automatically turns off the internal power supply 48 volt output when two or more external supplies are first powered on.

Example

The following are sample outputs of this command:

Slot-2 Stack.40 # show power

Slots	Type	Internal PSU	External PSU	External PSU	External PSU
Slot-1	X450e-24p	F	P		
Slot-2	X450a-24t	P	-		
Slot-3	X450a-24tDC	-	P		
Slot-4	X450a-48t	P	P		
Slot-5	X450a-24x	P	-		
Slot-6	X450a-24xDC	-	P		
Slot-7	X450e-48p	P	P	F	-
Slot-8	X450-24t	P	-		

Flags : (P) Power available, (F) Failed or no power,
 (O) 48V powered off when 2 or 3 external PSU are powered on,
 (-) Empty

Slot-2 Stack.41 #

Slot-2 Stack250.4 # show power

Slots	Type	Internal PSU	External PSU	External PSU	External PSU
Slot-1	X250e-24p	F	P		
Slot-2	X250a-24t	P	-		
Slot-3	X250-24xTAA	P	-		
Slot-4	X250a-48t	P	P		
Slot-5	X250a-24x	P	-		
Slot-6	X250e-48p	O	P	P	F
Slot-7					
Slot-8	X450-24x	P	-		

Flags : (P) Power available, (F) Failed or no power,
 (O) 48V powered off when 2 or 3 external PSUs are powered on,
 (-) Empty

Slot-2 Stack250.5 #

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms, but the output described in this section is available only on SummitStack.

show stacking

show stacking

Description

The show stacking command shows a summary of the nodes in the stack topology. The show stacking command will show all nodes that are in the stack topology.

The following information is displayed:

- Stack Topology is a ring or daisy-chain
- Active Topology is a ring or daisy-chain (or "This node is not in an Active Topology.")
- For each node:
 - node MAC address (factory assigned)
 - slot number in use
 - Stack State:
 - Disabled - Node is not configured for stacking.
 - Failed - Node can't come up in the stack because it has a duplicate slot number.
 - Overflow - The node has detected that there are more nodes in the stack topology than are allowed.
 - Listening - Initial state of attempt to join the stack
 - Stabilizing - Node is advertising its view of the topology and waiting until it sees no new topology changes.
 - Waiting - Topology has stabilized and stacking link blocking is being performed.
 - Active - The node is an active node and is fully programmed to operate in the active topology.
- Node role (master, backup, standby, or other transient node state).
- Flags describing the node's membership in an active topology
- Whether or not the node is "this" node, i.e., the node on which the command is run

Syntax Description

This command has no arguments or variables.

Default

There is not default value for this command.

Usage Guidelines

The asterisk (*) that precedes the node MAC address indicates the node on which this command is being executed, i.e., the node to which the user is logged in.

The node MAC address is the address that is factory assigned to the stackable.

The slot number shown is the number currently in use by the related node. Since slot number configuration only takes effect during node initialization, a change in configured value alone does not cause a change to the slot number that is in use. Slot numbers show as hyphen (-) characters on nodes that have stacking disabled.

The Stack State shows the state values.

The Role is one of the following: master, backup, standby, or <none>.

In a ring topology, the node on which this command is executed is always the first node displayed. The order of the nodes shown in the display is the order of their physical connection in the ring.

Even though the stack topology can be a ring, the active topology can simultaneously be a daisy chain because it can be only a proper subset of the stack topology. If the node on which this command is being executed is not active, the line

Active Topology is a ____

is replaced by the line

This node is not in an Active Topology.

The daisy chain topology is displayed in the order of physical connection. The master node detects the nodes that do not have stacking links to other members of the active topology, and these nodes become the ends of the stack. Such nodes will always display at the top and bottom of the output.

It is possible for a node to be in Stabilizing or Waiting state and still be in the active topology. This is because it is possible for a node that has become an active node to move to these states when a topology change is detected. Once having become an active node, the node remains an active node until it reboots or an overflow condition occurs.

The Flags have the following definitions:

- The "C" flag indicates that the related node is a candidate for membership of the same active topology to which the node on which the command is executed would belong.
- The "A" flag indicates that the related node is an active node in the active topology of which the node on which the command is run is also a candidate node. Being an active node is necessary but not sufficient for presence of the node in a slot. Once the node has fully initialized, the active node will appear as present in the `show slot` display.
- The "O" flag indicates that the related node is probably an active node in an active topology for which the node on which this command is being run is not a Candidate.

The "O" flag is useful for the case where there is an Inhibited link or a Disabled or Failed node that separates two active topologies. One active topology may contain the local node, and all other nodes in this active topology will not have the "O" flag set. All nodes that are members of an active topology that is separated by an Inhibited link from the active topology that contains the local node will have only the "O" flag set. All possibly active nodes will have the "O" flag set if the local node is not a member of any active topology. For any node for which the "O" flag is set, the "C" and "A" flags will not be set and vice-versa.

Example

The following example shows the output of show stacking command:

```
Slot-1 Stack.30 # show stacking
Stack Topology is a Ring
Active Topology is a Daisy-Chain
Node MAC Address      Slot  Stack State  Role      Flags
-----
*00:04:96:26:60:DD    1      Stabilizing  Master    CA-
00:04:96:26:60:EE    4      Stabilizing  Standby   C--
00:04:96:26:60:FF    -      Disabled    Master    ---
00:04:96:26:60:AA    -      Disabled    Master    ---
00:04:96:26:60:88    -      Disabled    Master    ---
00:04:96:26:60:99    -      Disabled    Master    ---
00:04:96:26:60:BB    2      Stabilizing  Standby   C--
00:04:96:26:60:CC    3      Active      Backup    CA-
(*) Indicates This Node
Flags: (C) Candidate for this active topology, (A) Active node,
       (O) node may be in Other active topology
Slot-1 Stack.31 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

show stacking configuration

```
show stacking configuration
```

Description

Shows how the nodes are configured in a stack topology. The configured values shown are the ones actually stored in the remote nodes at the time you issue this command.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

Some stacking parameters do not take effect until the next restart, the configured values and the values currently being used are both shown. Specifically, this applies to the slot number, whether or not stacking is enabled, master-capable configuration, the license level restriction, and stack MAC configuration.

The only parameters that take effect without a reboot are the node priority and the alternate management IP subnetwork and gateway.

The Stack MAC In Use has the following values:

- If the command is executed on the master node:
 - <none> if there is no stack MAC configured
 - The stack MAC configured on the master node otherwise
- If the command is executed on a non-master node:
 - <unknown>. The stack MAC address is only known by the executing master node. In this case, the M and m flags are not set. The i flag is set if there is a Stack MAC configured locally, and not otherwise.

Identified with the asterisk, the current node is the one on which the `show stacking configuration` command is executed.

A node identified with the "?" character indicates that timely attempts to fetch the configuration information from the node have failed. There are two possible reasons for this display:

- Communications with the node have been lost, in which case the node will probably be removed from the stack topology shortly.
- The node is too busy to respond in time.

A row that displays the "?" indicator will show the last values that were received from the node. If no values were ever received, all configured values will show as not configured (-) or <none>. The node MAC address and the slot number that is currently in use can still be displayed.

Example

The following example:

```
Slot-1 Stack.31 # show stacking configuration
```

```
Stack MAC in use: 02:04:96:26:60:DD
```

Node MAC Address	Slot Cfg Curr Pri	Alternate Mgmt IP / Mask	Alternate Gateway	Flags	Lic
*00:04:96:26:60:DD	1 1 1	192.168.130.101/24	192.168.130.001	CcEeMm-	Ee
00:04:96:26:60:EE	4 4 1	192.168.130.104/24	192.168.130.001	CcEeMm-	Ee
00:04:96:26:60:FF	6 - 1	192.168.130.106/24	192.168.130.001	---e-m-	--
?00:04:96:26:60:AA	7 - -	<none>	<none>	-c-----	--
00:04:96:26:60:88	8 - 1	192.168.130.108/24	192.168.130.001	---e--i	--
00:04:96:26:60:99	5 - -	<none>	<none>	-c---m-	--
00:04:96:26:60:BB	2 2 1	192.168.130.102/24	192.168.130.001	CcEeMm-	Ee
00:04:96:26:60:CC	3 3 1	192.168.130.103/24	192.168.130.001	CcEeMm-	Ee

* - Indicates this node

? - Cached information, node is not responding

Flags: (C) master-Capable in use, (c) master-capable is configured,
 (E) Stacking is currently Enabled, (e) Stacking is configured Enabled,
 (M) Stack MAC in use, (m) Stack MACs configured and in use are the same,
 (i) Stack MACs configured and in use are not the same or unknown,
 (-) Not in use or not configured

License Level Restrictions : (C) Core, (A) Advanced edge, or (E) Edge in use

(c) Core, (a) Advanced edge, or (e) Edge configured

(-) Not in use or not configured

```
Slot-1 Stack.32 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

show stacking detail

```
show stacking {node-address <node-address> | slot <slot-number>} detail
```

Description

This command displays information about a specified node.

Syntax Description

node-address	The MAC address of the node for which you want to see the stacking details.
slot-number	The slot number of the node in the stack. The value can be from 1 to 8.

Default

N/A

Usage Guidelines

If no node is specified, the output is generated for all nodes in the stack topology. If the specified node does not exist, an error message is displayed. The slot parameter is available only for active nodes in the same active topology as the node on which the command is being run. The node-address parameter is always available.

Current information represents stacking states and configured values that are currently in effect. *Configured* information is that which takes effect at node reboot only. Thus, differences between values in use and values configured can be seen here. The advantages of this command over the [show stacking configuration](#) command is that the values in use and configured values are fully expanded without the need for flags. You can also see the port state information of the node(s).

The roles values are: master, backup, standby, and None.

License level restrictions can be Edge, Advanced Edge, or Core.

Example

The following is a sample output of this command:

```
Slot-1 Stack.33 # show stacking slot 1 detail
Stacking Node 00:04:96:26:60:DD information:
  Current:
    Stacking                : Enabled
    Role                    : Master
    Priority                 : 2
    Slot number             : 1
    Stack state              : Active
    Master capable           : Yes
    License level restriction : Advanced edge
    In active topology?     : Yes
    Factory MAC address      : 00:04:96:26:60:DD
    Stack MAC address        : 02:04:96:26:60:DD
    Alternate IP address     : 192.168.130.101/24
    Alternate gateway        : 192.168.130.1
    Stack port 1:
      State                  : Operational
      Blocked?               : No
      Control path active?   : Yes
    Stack port 2:
      State                  : Operational
      Blocked?               : No
      Control path active?   : Yes
  Configured:
    Stacking                : Enabled
    Master capable           : Yes
    Slot number             : 1
    Stack MAC address        : 02:04:96:26:60:DD
    License level restriction : Edge
Slot-1 Stack.34 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

show stacking stack-ports

```
show stacking stack-ports
```

Description

The command displays the port states of each node in the stack topology and the connections between the nodes.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

The port and node MAC Address field values in a row of the output identify a particular stacking port. Each node MAC address appears twice in two consecutive rows in the output because each node has two stacking ports. On all platforms, the ports are labeled with the values 1 or 2. The order in which stacking ports appear in the display is the order in which they are physically connected. In the example below, note that the port numbers shown for the node with ID 00:04:96:26:60:CC are reversed. This means that port 2 of node 00:04:96:26:60:BB is connected to port 2 of node 00:04:96:26:60:CC. Similarly, port 2 of node 00:04:96:26:60:DD is connected to port 1 of node 00:04:96:26:60:EE, and since the topology is a ring, port 1 of node 00:04:96:26:60:CC is connected to port 1 of node 00:04:96:26:60:DD.

The Port State field contains one of the following states:

- Link Down – port is not receiving signal
- No Neighbor – the port is receiving signal but is not seeing a stack neighbor
- Overflow – 17 nodes (or more) are physically connected to this port.
- Inhibited – When you connected the link, active topologies were detected on both sides, and at least one slot number was duplicated. The stack merge is prevented.
- Operational – the port is operational in the stack. This is a necessary but insufficient condition for the port to be used for control path or user data. For example, a node with stacking Failed state may still show its port states as Operational.

The Flags field contains the following flag definitions:

- C - The control path is active on this port. Note that the user data path over the stack links follows the control path.
- B - The port is blocked from transmitting traffic that is to be flooded to multiple non-stacking ports. This flag is only set in an active ring topology. It will be set on two adjacent ports. In the example below, the active topology is a daisy chain, so no ports are blocked.

Identified with the asterisk, the current node is the one on which the “show stacking” command variant is executed. The stack topology is shown in a particular order. In a ring topology, the current node is

always the first node, the next node is the node connected to the port 2 of the first node, and the last node is the node connected to the port 1 of the current node. In a daisy chain, the order shown depends on the connection of the node on which the command executes:

- The first node is the one at the far end of the daisy-chain connected to the current node port 1.
- The last node is the one at the far end of the daisy-chain connected to the current node port 2.
- The previous node is the one at the near end of the daisy-chain connected to the current node port 1.
- The next node is the one at the near end of the daisy-chain connected to the current node port 2.
- If there is no node connected to the current node port 1, the current node is the first node.
- If there is no node connected to the current node port 2, the current node is the last node.

The slot number shown is the slot number in use on stacking enabled nodes. If the node does not have stacking enabled, a hyphen character ("-") is shown instead of a number.

Example

The following example shows the output for show stacking stack-ports:

```
Slot-1 Stack.32 # show stacking stack-ports
Stack Topology is a Ring
Slot Port Node MAC Address Port State Flags
----
*1 1 00:04:96:26:60:DD Operational C-
*1 2 00:04:96:26:60:DD Operational C-
4 1 00:04:96:26:60:EE Operational C-
4 2 00:04:96:26:60:EE Operational --
- 1 00:04:96:26:60:FF Operational --
- 2 00:04:96:26:60:FF Operational --
- 1 00:04:96:26:60:AA Operational --
- 2 00:04:96:26:60:AA Operational --
- 1 00:04:96:26:60:88 Operational --
- 2 00:04:96:26:60:88 Operational --
- 1 00:04:96:26:60:99 Operational --
- 2 00:04:96:26:60:99 Operational --
2 1 00:04:96:26:60:BB Operational --
2 2 00:04:96:26:60:BB Operational C-
3 2 00:04:96:26:60:CC Operational C-
3 1 00:04:96:26:60:CC Operational C-
(*) Indicates This Node
Flags: (C) Control path is active, (B) Port is Blocked
Slot-1 Stack.33 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

synchronize stacking

```
synchronize stacking {node-address <node-address> | slot <slot-number>}
```

Description

This command copies certain NVRAM based configuration parameters to the target node.

Syntax Description

node-address	The MAC address of the node. The node should be a valid node in the stack topology.
slot-number	The number of a slot that is associated to a node that is present in the active topology.

Default

N/A

Usage Guidelines

This command synchronizes the following NVRAM-based parameters:

- stacking mode
- stack MAC address
- failsafe account and password
- failsafe account access point permissions (whether the failsafe account is allowed over the stacking links, console port, or management port)
- the selected partition

These parameters are copied from the executing node's NVRAM to the target node's NVRAM.

Example

Example for the synchronize stacking command output:

```
Slot-2 Stack.3 > synchronize stacking slot 3
Are you sure you want to synchronize the specified slot with this slot's stacking
configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified
node(s).
Slot-2 Stack.4 >

Slot-2 Stack.4 > synchronize stacking node 00:04:96:27:87:10
Are you sure you want to synchronize the specified node with this node's stacking
configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified
node(s).
Slot-2 Stack.5 >

Slot-2 Stack.5 > synchronize stacking
Are you sure you want to synchronize all remote nodes with this node's stacking
configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified
node(s).
Slot-2 Stack.6 >
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

unconfigure stacking

```
unconfigure stacking {node-address <node-address> | slot <slot-number>}
```

Description

This command resets all stacking parameters to the default or unconfigured values shown in [Table 24](#). If a node-address or slot-number is not supplied, the command applies to all nodes in the stack topology.

Table 24: Stacking Configuration Items, Time of Effect and Default Value

Configuration Item	Takes Effect	Default Value
Stacking Mode	At boot time	Disabled
Slot Number	At boot time	1
Master-Capable	At boot time	Yes
License Restriction	At boot time	Not Configured
Priority	At next master election	Automatic
Alternate IP Address	Immediately	Not Configured
Stack MAC	At boot time	Not Configured

Syntax Description

node-address	The MAC address of the node on which you want stacking parameters unconfigured.
slot-number	The number of a slot that is associated to a node that is present in the active topology.

Default

N/A

Usage Guidelines

Run this command from any node.

Example

To unconfigure the stacking parameters of all nodes in the stack topology:

```
unconfigure stacking
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and Summit Family switches.

unconfigure stacking alternate-ip-address

```
unconfigure stacking {node-address <node-address> | slot <slot-number>}  
alternate-ip-address
```

Description

Removes the configured alternate management IP address from the specified node. If no node is specified, the alternate management IP address is removed from every node. The change takes effect immediately for all nodes operating in stacking mode.

Syntax Description

node-address	The MAC address of the node on which you want to unconfigure stacking alternate IP address.
slot-number	The slot number of a node in the active topology.

Default

N/A

Usage Guidelines

Run this command from any node.

Example

To unconfigure stacking alternate-ip-address on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed alternate-ip-address
```

To unconfigure stacking alternate-ip-address configured on the active node in slot 4:

```
unconfigure stacking slot 4 alternate-ip-address
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

unconfigure stacking license-level

```
unconfigure stacking {node-address <node-address> | slot <slot-number>}  
license-level
```

Description

This command removes a previously configured license level restriction.

Syntax Description

node-address	The MAC address of the node on which you want to unconfigure stacking license level.
slot-number	The slot number of a node in the active topology.

Default

N/A

Usage Guidelines

If no node is specified, the licensing restriction is removed from all nodes in the stack topology.

After the command is executed, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To unconfigure stacking license level on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed license-level
```

To unconfigure stacking license level configured on slot 4:

```
unconfigure stacking slot 4 license-level
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack and the Summit family switches.

23 EAPS Commands

This chapter describes commands for configuring and monitoring Ethernet Automatic Protection Switching (EAPS).

The EAPS protocol provides fast protection switching to layer 2 switches interconnected in an Ethernet ring topology, such as a metropolitan area network (MAN) or large campuses. EAPS protection switching is similar to what can be achieved with the Spanning Tree Protocol (STP), but offers the advantage of converging in less than a second when a link in the ring breaks.

EAPS operates by declaring an EAPS domain on a single ring. Any VLAN that warrants fault protection is configured on all ring ports in the ring, and is then assigned to an EAPS domain. On that ring domain, one switch, or node, is designated the *master* node, while all other nodes are designated as *transit* nodes.

One port of the master node is designated the master node's *primary* port (P) to the ring; another port is designated as the master node's *secondary* port (S) to the ring. In normal operation, the master node blocks the secondary port for all non-control traffic belonging to this EAPS domain. If the master node detects a break in the ring, it unblocks its secondary port and allows data traffic to be transmitted and received through it.

EAPS fault detection on a ring is based on a single *control* VLAN per EAPS domain. This EAPS domain provides protection to one or more data-carrying VLANs called *protected* VLANs. The control VLAN is used only to send and receive EAPS messages; the protected VLANs carry the actual data traffic. As long as the ring is complete, the EAPS master node blocks the protected VLANs from accessing its secondary port.

A master node detects a ring fault in any of three ways:

- "Link down" message sent by a transit node on the control VLAN
- Ring port down event from lower hardware layers
- Failed response to a periodic health-check packet on the control VLAN

When the master node detects a failure, it declares a "failed" state and opens its logically blocked secondary port on all the protected VLANs. The master node also flushes its forwarding database (FDB) and sends a message on the control VLAN to all of its associated transit nodes to flush their forwarding databases.

To take advantage of the Spatial Reuse technology and broaden the use of the ring's bandwidth, EAPS supports multiple EAPS domains running on the ring at the same time.

Licensing

You must have a Core or an Advanced Core license to configure and use all of the EAPS features described in this chapter.

A subset of EAPS, called EAPS Edgemode, is available with an Edge or an Advanced Edge license. The following features are available with EAPS Edgemode:

- Switches can belong to only one EAPS ring.
- You can have up to four EAPS domains using two matching ring ports.

EAPS shared ports is not supported with an Edge or an Advanced Edge license.

The BlackDiamond 10808 switch with an MSM-1 module or an MSM1-XL module ships with a Core or an Advanced Core license, respectively.

The BlackDiamond 12804 switch with an MSM-5 module or an MSM-5R module ships with a Core license. The BlackDiamond 12802 switch ships with an Advanced Edge license.

The BlackDiamond 8800 series switch with an MSM-G8X module or an MSM-48 module, Summit X450 series switch, and Summit X450a series switch ship with an Advanced Edge license. To use the complete EAPS functionality, including running two or more EAPS rings, having a switch belong to multiple EAPS rings, or configuring shared-ports that allow multiple EAPS domains to share a common link, you must upgrade to a Core software license.

A SummitStack with a core license provides all of the EAPS features described in this chapter.

The Summit X450e series switches ship with an Edge license. You can upgrade to an Advanced Edge license; however, both the Edge and the Advanced Edge licenses only support EAPS Edgemode. You cannot upgrade to a Core software license. You cannot use the Summit X450e platform in a SummitStack when the SummitStack is operating with a Core license level.

**NOTE**

The Summit X450e series switch only supports EAPS Edgemode.

For more information about software licensing, including how to obtain and upgrade your license, see [Chapter 2, “ExtremeXOS Overview,”](#) in the *ExtremeXOS Concepts Guide*.

clear eaps counters

```
clear eaps counters
```

Description

Clears, resets the counters gathered by EAPS for all of the EAPS domains and any EAPS shared ports configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to clear, reset the EAPS counters.

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

To display information about the EAPS counters, use the following commands:

- `show eaps counters`—This command displays summary EAPS counter information.
- `show eaps counters shared-port`—If configured for EAPS shared ports, this command displays summary EAPS shared port counter information.

Example

The following command clears, resets all of the counters for EAPS:

```
clear eaps counters
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure eaps add control vlan

```
configure eaps <name> add control vlan <vlan_name>
```

Description

Adds the specified control VLAN to the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

You must configure one *control* VLAN for each EAPS domain. The control VLAN is used only to send and receive EAPS messages.

The VLAN that will act as the control VLAN must be configured as follows:

- The VLAN must NOT be assigned an IP address, to avoid loops in the network.
- Only ring ports may be added as members of the control VLAN.
- The ring ports of the control VLAN must be tagged.

A control VLAN cannot belong to more than one EAPS domain. When the EAPS domain is active, you cannot delete or modify the configuration of the control VLAN.

By default, EAPS protocol data units (PDUs) are automatically assigned an 802.1p priority of seven. With a priority of seven, EAPS PDUs take precedence within the EAPS topology thereby giving priority to EAPS traffic, including EAPS control VLAN traffic and control VLAN messages. This ensures that the control VLAN messages reach their intended destinations. You do not need to configure a QoS profile for the control VLAN.

However, if you have a protected (data) VLAN with a QoS profile of QP8, you may need to create a QoS profile of QP8 for the control VLAN. Whether you need to configure the control VLAN with a QoS profile of QP8 depends entirely on your configuration.



NOTE

On the BlackDiamond 8800 original modules and the original Summit X450 series switch (whether or not included in a SummitStack), assigning a QoS profile to the control VLAN affects the amount of available ACL masks. For more information see [“Conserving ACL Masks and Rules—BlackDiamond 8800 Original Series Modules and the Summit X450 Series Switches Only”](#) in the ExtremeXOS Concepts Guide.

The VLAN must already exist before you can add it as a control VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.8 # configure eaps megtest add control foo
                                         ^
%% Invalid input detected at '^' marker.
```

To create the VLAN, use the `create vlan` command.

Example

The following command adds the control VLAN “keys” to the EAPS domain “eaps_1.”

```
configure eaps eaps_1 add control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps add protected vlan

```
configure eaps <name> add protected vlan <vlan_name>
```

Description

Adds the specified protected VLAN to the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the protected VLAN.

Default

N/A.

Usage Guidelines

You must configure one or more *protected* VLANs for each EAPS domain. The protected VLANs are the data-carrying VLANs.

A protected VLAN can be added to one or more EAPS domains.

When you configure the VLAN that will act as a protected VLAN, the ring ports of the protected VLAN must be tagged (except in the case of the default VLAN). As long as the ring is complete, the master node blocks the protected VLANs on its secondary port.



NOTE

If the protected VLAN has a QoS profile of QP8, make sure the control VLAN has the appropriate QoS profile setting. For more information, see the command [configure eaps add control vlan](#) on page 1358.

The VLAN must already exist before you can add it as a protected VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.5 # configure eaps megtest add protected foo
                                     ^
%% Invalid input detected at '^' marker.
```

To create the VLAN, use the [create vlan](#) command.

Example

The following command adds the protected VLAN “orchid” to the EAPS domain “eaps_1”:

```
configure eaps eaps_1 add protected vlan orchid
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps config-warnings off

```
configure eaps config-warnings off
```

Description

Disables the loop protection warning messages displayed when configuring specific EAPS parameters.

Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.

Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a VLAN
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

Extreme Networks recommends that you keep the loop protection warning messages enabled. If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For example, if you use a script to configure your EAPS settings, disabling the warning messages allows you to configure EAPS without replying to each interactive yes/no question.

To confirm the setting on the switch, use the `show eaps {<eapsDomain>} {detail}` command.

Example

The following command disables the loop protection warning messages:

```
configure eaps config-warnings off
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps config-warnings on

```
configure eaps config-warnings on
```

Description

Enables the loop protection warning messages displayed when configuring specific EAPS parameters.

Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.

Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a VLAN
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

Extreme Networks recommends that you keep the loop protection warning messages enabled.

Example

The following command enables the loop protection warning messages:

```
configure eaps config-warnings on
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps delete control vlan

```
configure eaps <name> delete control vlan <vlan_name>
```

Description

Deletes the specified control VLAN from the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the control VLAN “keys” from the EAPS domain “eaps_1”:

```
configure eaps eaps_1 delete control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps delete protected vlan

```
configure eaps <name> delete protected vlan <vlan_name>
```

Description

Deletes the specified protected VLAN from the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the protected VLAN.

Default

N/A.

Usage Guidelines

To prevent loops in the network, you must delete the ring ports (the primary and the secondary ports) from the protected VLAN **before** deleting the protected VLAN from the EAPS domain. Failure to do so may cause a loop in the network.

The switch displays by default a warning message and prompts you to delete the VLAN from the EAPS domain. When prompted, do one of the following:

- Enter **y** delete the VLAN from the specified EAPS domain.
- Enter **n** or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command [on page 1362](#).

Useful show Commands. Use the following `show` commands to display information about your EAPS domain, including protected VLANs and primary and secondary ports:

- `show vlan`—This command displays summary information for all of the VLANs on the device. If the VLAN is a protected VLAN, the P flag appears in the flag column. To see more detailed information about the protected VLAN, use the `show vlan <vlan_name> detail` command.
- `show eaps`—This command displays summary EAPS domain information, including the name of the domain and the primary and secondary ports. To see more detailed information, including the name of the protected VLAN and the primary and secondary ports, use the `show eaps <eapsDomain>` command.
- `show vlan eaps`—This command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

Example

The following command deletes the protected VLAN “orchid” from the EAPS domain “eaps_1”:

```
configure eaps eaps_1 delete protected vlan orchid
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure EAPS ring-ports are deleted from the VLAN first.  Otherwise
         deleting the VLAN from the EAPS domain could cause a loop in the network!
Are you sure you want to remove the VLAN before deleting EAPS ring-ports.? (y/n)
```

Enter *y* to delete the VLAN from the specified EAPS domain. Enter *n* to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps failtime

```
configure eaps <name> failtime <seconds>
```

Description

Configures the value of the failtimer the master node uses for EAPS health-check packets.

Syntax Description

name	Specifies the name of an EAPS domain.
seconds	Specifies the number of seconds the master node waits to receive a health-check packet before the failtimer expires. Default is 3 seconds, and the range is 2 to 300 seconds.

Default

The default is 3 seconds.

Usage Guidelines

Use the `failtime` keyword and its associated `seconds` parameter to specify the amount of time the master node waits before the failtimer expires. The `seconds` parameter must be set greater than the configured value for `hellotime`. The default value is three seconds.

Increasing the `failtime` value provides more protection by waiting longer to receive a health-check packet when the network is congested.



NOTE

You configure the action taken when the failtimer expires by using the `configure eaps failtime expiry-action` command.

In ExtremeXOS 11.0, the failtimer range was 2 to 60 seconds.

Example

The following command configures the failtimer value for the EAPS domain "eaps_1" to 15 seconds:

```
configure eaps eaps_1 failtime 15
```

History

This command was first available in ExtremeXOS 11.0.

The range for the failtimer was changed to 2 to 300 seconds in ExtremeXOS 11.1. The default value for the failtimer remains unchanged.

Platform Availability

This command is available on all platforms.

configure eaps failtime expiry-action

```
configure eaps <name> failtime expiry-action [open-secondary-port | send-alert]
```

Description

Configures the action taken when the failtimer expires.

Syntax Description

name	Specifies the name of an EAPS domain.
open-secondary-port	Specifies to open the secondary port when the failtimer expires.
send-alert	Specifies that a critical message is sent to the syslog when the failtimer expires.

Default

Default is send-alert.

Usage Guidelines

By default the action is to send an alert if the failtimer expires. Instead of going into a “Failed” state, the master node remains in a “Complete” or “Init” state, maintains the secondary port blocking, and writes a critical error message to syslog warning the user that there is a fault in the ring. An SNMP trap is also sent.

If the EAPS ring contains non-EAPS devices, you must use the `open-secondary-port` parameter.



NOTE

Use caution when setting the failtimer expiry action to `open-secondary port`. Using this configuration, if the master node loses three consecutive hello PDUs, the failtimer will expire—but there may not be a break in the ring. Opening the secondary port in this situation would create a loop.

Example

The following command configures the failtimer expiry action for EAPS domain “eaps_1”:

```
configure eaps eaps_1 failtime expiry-action open-secondary-port
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps fast-convergence

```
configure eaps fast-convergence [off | on]
```

Description

Enables EAPS to converge more quickly.

Syntax Description

off	Turns fast-convergence off. Default is off.
on	Turns fast-convergence on.

Default

Default is off.

Usage Guidelines

This command acts on the switch, not per domain.

In certain environments to keep packet loss to a minimum when the ring is broken, configure EAPS with fast-convergence turned on. If fast convergence is turned on, it will be displayed under the [show eaps](#) command. For example:

```
show eaps
```

```
EAPS Enabled: No
EAPS Fast-Convergence: Off
Number of EAPS instances: 0
# EAPS domain configuration :
```

Domain	State	Mo	En	Pri	Sec	Control-Vlan	VID	Count



NOTE

If fast-convergence is turned on, the link filters on all EAPS ring ports are turned off. This can result problems if the port's hardware encountered a problem and started "flapping" between link-up/link-down states.

Example

The following command configures fast convergence for all of the EAPS domains on the switch:

```
configure eaps fast-convergence on
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps hellotime

```
configure eaps <name> hellotime <seconds>
```

Description

Configures the value of the hello timer the master node uses for the EAPS health-check packet.

Syntax Description

name	Specifies the name of an EAPS domain.
seconds	Specifies the number of seconds to wait between transmission of the health-check packets on the control VLAN. Must be greater than 0; default is 1 second, and the range is 1 to 15 seconds.

Default

Default is 1 second.

Usage Guidelines

Use the `hellotime` keyword and its associated `seconds` parameter to specify the amount of time the master node waits between transmissions of health-check packets on the control VLAN. Increasing the `hellotime` value keeps the processor from sending and processing too many health-check packets.

This command applies only to the master node. If you configure the polling timers for a transit node, those values are ignored. If you later reconfigure that transit node as the master node, the polling timer values is used as the current values.

In ExtremeXOS 11.0, the range is 1 to 15 seconds. If you are running ExtremeXOS 11.0 with the hello timer value greater than 15 seconds and you upgrade to ExtremeXOS 11.1 or later, you must modify the hello timer to be within the 1 to 15 seconds range.

Example

The following command configures the hellotime value for the EAPS domain “eaps_1” to 2 seconds:

```
configure eaps eaps_1 hellotime 2
```

History

This command was first available in ExtremeXOS 11.0.

The range for the hello timer was changed to 1 to 15 seconds in ExtremeXOS 11.1. The default value for the hello timer remains unchanged.

Platform Availability

This command is available on all platforms.

configure eaps mode

```
configure eaps <name> mode [master | transit]
```

Description

Configures the switch as either the EAPS master node or as an EAPS transit node for the specified domain.

Syntax Description

name	Specifies the name of an EAPS domain.
master	Specifies that this switch should be the master node for the named EAPS domain.
transit	Specifies that this switch should be the transit node for the named EAPS domain.

Default

N/A.

Usage Guidelines

One node (or switch) on the ring *must* be configured as the master node for the specified domain; all other nodes (or switches) on the ring are configured as transit nodes for the same domain.

If you configure a switch to be a transit node for an EAPS domain, the switch displays by default messages to:

- Remind you to configure a master node in the EAPS domain.
- Notify you that changing a master node to a transit node might cause a loop in the network. If you have not assigned a new master node before changing the current master node to a transit node, you might cause a loop in the network.

When prompted, do one of the following:

- Enter `y` to identify the switch as a transit node.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command on page 1362.

Example

The following command identifies this switch as the master node for the domain named `eaps_1`:

```
configure eaps eaps_1 mode master
```

The following command identifies this switch as a transit node for the domain named eaps_1:

```
configure eaps eaps_1 mode transit
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure this specific EAPS domain has a Master node in the ring.  If
        you change this node from EAPS master to EAPS transit, you could cause a
        loop in the network.
```

```
Are you sure you want to change mode to transit? (y/n)
```

Enter *y* to identify the switch as a transit node. Enter *n* to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps name

```
configure eaps <old_name> name <new_name>
```

Description

Renames an existing EAPS domain.

Syntax Description

old_name	Specifies the current name of an EAPS domain.
new_name	Specifies a new name for the EAPS domain.

Default

N/A.

Usage Guidelines

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command renames EAPS domain “eaps-1” to “eaps-5”:

```
configure eaps eaps-1 name eaps-5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps port

```
configure eaps <name> [primary | secondary] port <ports>
```

Description

Configures a node port as the primary or secondary port for the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
primary	Specifies that the port is to be configured as the primary port.
secondary	Specifies that the port is to be configured as the secondary port.
ports	Specifies one port or slot and port.

Default

N/A.

Usage Guidelines

Each node on the ring connects through two ring ports. One port must be configured as the *primary* port; the other must be configured as the *secondary* port.

The primary and secondary ports have significance only on a master node. The health-check messages are sent out the primary port of the master node, and the master node blocks the protected VLANs on the secondary port.

You cannot use the same port for the master node's secondary port and the shared-port. The master node's secondary port must be different from the shared-port.

There is no distinction between the primary and secondary ports on a transit node.

Beginning with ExtremeXOS 11.1, if you have a primary or secondary port that is a member of a load-shared group, you do not need to disable your EAPS domain and remove that ring port when modifying the load-shared group. For more information about configuring load sharing on your switch, see [Chapter 6, "Configuring Slots and Ports on a Switch,"](#) in the *ExtremeXOS Concepts Guide*.

You must have a Core or an Advanced Core license to configure multiple EAPS domains with different ring ports.

Messages Displayed when Adding EAPS Ring Ports to a VLAN. If you attempt to add EAPS ring ports to a VLAN that is not protected by EAPS, the switch prompts you by default to confirm this action. For example, if you use the `configure vlan <vlan_name> add ports <port_list>` command, and the ports that you are attempting to add to the VLAN are currently used by EAPS as either primary or secondary ring ports, the switch displays the following message:

```
Make sure <vlan_name> is protected by EAPS. Adding EAPS ring ports to a VLAN could
cause a loop in the network.
```

```
Do you really want to add these ports (y/n)
```

Enter `y` to add the ports to the VLAN. Enter `n` or press [Return] to cancel this action.

If you see this message, either configure the VLAN as an EAPS protected VLAN by using the `configure eaps add protected vlan` command or add ports that the EAPS domain does not use as primary or secondary ring ports.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off` command [on page 1362](#).

Example

The following command adds port 1 of the module installed in slot 8 to the EAPS domain “eaps_1” as the primary port:

```
configure eaps eaps_1 primary port 8:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps shared-port link-id

```
configure eaps shared-port <ports> link-id <id>
```

Description

Configures the link ID of the shared port.

Syntax Description

ports	Specifies the port number of the common link port.
id	Specifies the link ID of the port. The link ID range is 1 to 65535.

Default

N/A.

Usage Guidelines

Each common link in the EAPS network must have a unique link ID. The controller and partner shared ports belonging to the same common link must have *matching* link IDs. No other instance in the network should have that link ID.

If you have multiple adjacent common links, Extreme Networks recommends that you configure the link IDs in ascending order of adjacency. For example, if you have an EAPS configuration with three adjacent common links, moving from left to right of the topology, configure the link IDs from the lowest to the highest value.

Example

The following command configures the EAPS shared port 1:1 to have a link ID of 1.

```
configure eaps shared-port 1:1 link-id 1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

configure eaps shared-port mode

```
configure eaps shared-port <ports> mode <controller | partner>
```

Description

Configures the mode of the shared port.

Syntax Description

ports	Specifies the port number of the shared port.
controller	Specifies the controller mode. The controller is the end of the common link responsible for blocking ports when the common link fails thereby preventing the superloop.
partner	Specifies partner mode. The partner is responsible only for sending and receiving health-check messages.

Default

N/A.

Usage Guidelines

The shared port on one end of the common link must be configured to be the *controller*. This is the end responsible for blocking ports when the common link fails thereby preventing the superloop.

The shared port on the other end of the common link must be configured to be the *partner*. This end does not participate in any form of blocking. It is responsible only for sending and receiving health-check messages.

Example

The following command configures the shared port 1:1 to be the controller.

```
configure eaps shared-port 1:1 mode controller
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

configure eaps shared-port segment-timeout expiry-action

```
configure eaps shared-port <ports> segment-timeout expiry-action [segment-down | send-alert]
```

Description

Configures the action taken when a segment is down.

Syntax Description

ports	Specifies a partner or controller shared port.
segment-down	Specifies the behavior of the partner or controller switch. If either switch does not receive health check messages during the defined segment timeout, the switch marks the segment as down.
send-alert	Specifies the behavior of the partner or controller switch. If either switch does not receive health check messages during the defined segment timeout, the switch keeps the segment in the up state, with its failed flag set, and sends a warning message to the log.

Default

Default is send-alert.

Usage Guidelines

By default, the action is to send an alert if the failtimer expires. Instead of the segment going into a failed state and being marked as segment down, the segment remains in a segment up state with the failed flag set. The switch writes a critical error message to the syslog warning the user that there is a fault in the segment. An SNMP trap is also sent.

If the EAPS segment contains non-EAPS devices, you must use the `segment-down` parameter.



NOTE

Use caution when setting the segment-timeout expiry action to `segment-down`. Using this configuration, if the controller or partner node loses three consecutive hello PDUs, the failtimer will expire—but there may not be a break in the segment. Opening a blocked port in this situation would create a loop.

The following describes some general recommendations for using this command:

- When you configure your Extreme Networks switches as the partner and controller, respectively, make sure that their segment timer configurations are identical.

For example, if you have a partner switch with the segment-timeout expiry action set to `send-alert`, make sure the controller switch has its segment-timeout expiry action set to `send-alert`.

However, if you have a partner switch with the segment-timeout expiry action set to `send-alert`, and the controller switch does not have a segment timer configuration, you must configure the partner switch's segment-timeout expiry action to `segment-down`.

- If you have a network containing non-Extreme Networks switches, set the segment-timeout expiry action to `segment-down`.

The following events may cause a ring segment failure:

- There is a hardware failure.
- The controller or partner received a Link Down message from the partner or controller, respectively.
- The segment timer expires and the expiry action was set to segment down. This means that either the controller or partner did not receive health check messages during the defined segment timeout.

To view shared-port information, including shared-port segment status, use the `show eaps shared-port {<port>} {detail}` command.

Example

The following example makes the following assumptions:

- The controller switch is an Extreme Networks switch running ExtremeXOS 11.1
- The partner switch is an Extreme Networks switch running ExtremeWare 7.3

The following command configures the controller to keep the segment down if the segment-timeout expires:

```
configure eaps shared-port 3:3 segment-timeout expiry-action segment-down
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

create eaps

```
create eaps <name>
```

Description

Creates an EAPS domain with the specified name.

Syntax Description

name	Specifies the name of an EAPS domain to be created. May be up to 32 characters in length.

Default

N/A.

Usage Guidelines

The `name` parameter is a character string of up to 32 characters that identifies the EAPS domain to be created.

If you use the same name across categories (for example, STPD and EAPS names), Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command creates EAPS domain `eaps_1`:

```
create eaps eaps_1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create eaps shared-port

```
create eaps shared-port <ports>
```

Description

Creates an EAPS shared port on the switch.

Syntax Description

ports	Specifies the port number of the common link port.

Default

N/A.

Usage Guidelines

To configure a common link, you must create a shared port on each switch on either end of the common link.

Example

The following command creates a shared port on the EAPS domain.

```
create eaps shared-port 1:2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

delete eaps

```
delete eaps <name>
```

Description

Deletes the EAPS domain with the specified name.

Syntax Description

name	Specifies the name of an EAPS domain to be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes EAPS domain eaps_1:

```
delete eaps eaps_1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

delete eaps shared-port

```
delete eaps shared-port <ports>
```

Description

Deletes an EAPS shared port on a switch.

Syntax Description

ports	Specifies the port number of the Common Link port.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes shared port 1:1.

```
delete eaps shared-port 1:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

disable eaps

```
disable eaps {<name>}
```

Description

Disables the EAPS function for a named domain or for an entire switch.

Syntax Description

name	Specifies the name of an EAPS domain.

Default

Disabled for the entire switch.

Usage Guidelines

To prevent loops in the network, the switch displays by default a warning message and prompts you to disable EAPS for a specific domain or the entire switch. When prompted, do one of the following:

- Enter `y` to disable EAPS for a specific domain or the entire switch.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command [on page 1362](#).

BlackDiamond 8800 Series Switch, SummitStack, and Summit Family of Switches Only. If you configure EAPS spatial reuse and disable the EAPS domain on a transit node, this may cause the master node's fail timer to expire and EAPS PDUs to no longer be forwarded through out the domain. If this occurs, make sure the EAPS domain is enabled on all of the transit nodes. To confirm the current state of EAPS on the switch, use the [show eaps](#) command.

Example

The following command disables the EAPS function for entire switch:

```
disable eaps
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Disabling EAPS on the switch could cause a loop in the network!
Are you sure you want to disable EAPS? (y/n)
```

Enter `y` to disable EAPS on the switch. Enter `n` to cancel this action.

The following command disables the EAPS function for the domain "eaps-1":

```
disable eaps eaps-1
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Disabling specific EAPS domain could cause a loop in the network!  
Are you sure you want to disable this specific EAPS domain? (y/n)
```

Enter `y` to disable the EAPS function for the specified domain. Enter `n` to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable eaps

```
enable eaps {<name>}
```

Description

Enables the EAPS function for a named domain or for an entire switch.

Syntax Description

name	Specifies the name of an EAPS domain.

Default

Disabled.

Default command enables EAPS for the entire switch.

Usage Guidelines



NOTE

If you use the same name across categories (for example, STPD and EAPS names), you must specify the identifying keyword as well as the actual name.

To configure and enable an EAPS, complete the following steps:

- 1 Create EAPS domain and assign the name.
- 2 Configure the control VLAN.
- 3 Configure the protected VLAN(s).
- 4 Add the control VLAN to EAPS domain.
- 5 Add the protected VLAN(s) to EAPS domain.
- 6 Configure EAPS mode, master or transit.
- 7 Configure EAPS port, secondary and primary.
- 8 If desired, configure timeout and action for failtimer expiration*.
- 9 If desired, configure the hello time for the health-check packets*.
- 10 Enable EAPS for the entire switch.
- 11 If desired, enable Fast Convergence*.
- 12 Enable EAPS for the specified domain.

Although you can enable EAPS prior to configuring these steps, the EAPS domain(s) will not run until you configure these parameters. (The steps with * can be configured at any time, even after the EAPS domains are running.)

You must enable EAPS globally and specifically for each named EAPS domain.

BlackDiamond 8800 Series Switch, SummitStack, and Summit Family of Switches Only. If you configure EAPS spatial reuse and disable the EAPS domain on a transit node, this may cause EAPS PDUs from being forwarded through out the domain thereby causing the master node's fail timer to expire. If this occurs, make sure the EAPS domain is enabled on all of the transit nodes. To confirm the current state of EAPS on the switch, use the `show eaps` command.

Example

The following command enables the EAPS function for entire switch:

```
enable eaps
```

The following command enables the EAPS function for the domain "eaps-1":

```
enable eaps eaps-1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show eaps

```
show eaps {<eapsDomain>} {detail}
```

Description

Displays EAPS status information.

Syntax Description

eapsDomain	Specifies the name of an EAPS domain.
detail	Specifies all available detail for each domain.

Default

N/A.

Usage Guidelines

If you enter the `show eaps` command without a keyword, the command displays less than with the `detail` keyword.

Use the optional domain `eapsDomain` parameter to display status information for a specific EAPS domain.

Some state values are different on a transit node than on a master node.

The switch displays the following fields:

EAPS Enabled:	Current state of EAPS on this switch: • Yes—EAPS is enabled on the switch. • No—EAPS is not enabled.
EAPS Fast Convergence:	Displays only when Fast Convergence is on.
EAPS Display Config Warnings:	Displays the setting for loop protection messages: • On—Loop protection messages are displayed (this is the default behavior). • Off—Loop protection messages are not displayed.
Number of EAPS instances:	Number of EAPS domains created. The maximum number of EAPS domains per switch is 128.
Name:	The configured name for this EAPS domain.

State:	On a transit node, the command displays one of the following states: • Idle—The EAPS domain has been enabled, but the configuration is not complete. • Links-Up—This EAPS domain is running, and both its ports are up and in the FORWARDING state. • Links-Down—This EAPS domain is running, but one or both of its ports are down. • Preforwarding—This EAPS domain is running, and both of its ports are up, but one of them is in a temporary BLOCKED state. On a master node, the command displays one of the following states: • Idle—The EAPS domain has been enabled, but the configuration is not complete. • Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state. • Complete—The ring is in the COMPLETE state for this EAPS domain. • Failed—There is a break in the ring for this EAPS domain. • Pre-Init—The EAPS domain has started operation for Init state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from hardware layer indicating the operation is completed. • Pre-Complete—The EAPS domain has started operation for Complete state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from the hardware layer indicating the operation is completed. • [Failtimer Expired]—When the failtimer expires and it's action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node will continue to remain in COMPLETE or INIT state with it's secondary port blocking.
[Running: ...]	• Yes—This EAPS domain is running. • No—This EAPS domain is not running.
Enabled:	Indicates whether EAPS is enabled on this domain. • Y—EAPS is enabled on this domain. • N—EAPS is not enabled.
Mode:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Port status:	• Unknown—This EAPS domain is not running, so the port status has not yet been determined. • Up—The port is up and is forwarding data. • Down—The port is down. • Blocked—The port is up, but data is blocked from being forwarded.
Tag status:	Tagged status of the control VLAN: • Tagged—The control VLAN has this port assigned to it, and the port is tagged in the VLAN. • Untagged—The control VLAN has this port assigned to it, but the port is untagged in the control VLAN. • Undetermined—Either a VLAN has not been added as the control VLAN to this EAPS domain or this port has not been added to the control VLAN.

Hello Timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits between transmissions of health check packets.
Fail Timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits before the failtimer expires.
Failtimer expiry action:	Displays the action taken when the failtimer expires: • Send-alert—Sends a critical message to the syslog when the failtimer expires. • Open-secondary-port—Opens the secondary port when the failtimer expires. Displays only for master nodes.
Preforwarding Timer interval: ^a	The configured value of the timer. This value is set internally by the EAPS software. The set value is 15 seconds. Note: If two links in an EAPS domain go down at the same time and one link comes back up, it will take 15 seconds for the reconnected link to start receiving traffic again.
Last update:	Indicates the last time a hello packet is received from the master node.
EAPS Domain has ... Controller Vlan:	Lists the assigned name and ID of the control VLAN.
EAPS Domain has ... Protected Vlan: ^b	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain.
Number of Protected Vlan:	The count of protected VLANs configured on this EAPS domain.

a. These fields apply only to transit nodes; they are not displayed for a master node.

b. This list is displayed when you use the `detail` keyword in the `show eaps` command.

Example

The following display shows sample output for the command `show eaps`:

```
EAPS Enabled: Yes
EAPS Fast-Convergence: Off
EAPS Display Config Warnings: On
Number of EAPS instances: 2
# EAPS domain configuration :
-----
Domain          State           Mo  En  Pri  Sec  Control-Vlan VID  Count
-----
d1              Complete       M   Y   3:8  3:16  c1                (1000) 100
d2              Links-Up       T   Y   3:8  3:16  c2                (1001) 100
-----
```

The following display shows sample output for the command `show eaps <eapsDomain>`:

```
Name: d1
State: Complete                      Running: Yes
Enabled: Yes      Mode: Master
Primary port:    3:8      Port status: Up      Tag status: Tagged
Secondary port: 3:16     Port status: Blocked  Tag status: Tagged
Hello timer interval: 1 sec
Fail timer interval: 3 sec
Fail Timer expiry action: Send alert
Last update: From Master Id 00:01:30:f9:9c:b0, at Wed Jun  9 09:09:35 2004
EAPS Domain has following Controller Vlan:
Vlan Name      VID
c1             1000
```

EAPS Domain has following Protected Vlan(s):

Vlan Name	VID
p_1	1
p_2	2
p_3	3
p_4	4
p_5	5
p_6	6
p_7	7
p_8	8
p_9	9
p_10	10
p_11	11
p_12	12
p_13	13
p_14	14
p_15	15
p_16	16
p_17	17
p_18	18
p_19	19
p_20	20
p_21	21
p_22	22
p_23	23
p_24	24
p_25	25
p_26	26
p_27	27
p_28	28
p_29	29
p_30	30



NOTE

You may see a slightly different display, depending on whether you display the master node or the transit node.

The display from the `show eaps detail` command shows all the information shown in the `show eaps <eapsDomain>` command, but displays information for all configured EAPS domains.

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show eaps counters

```
show eaps counters [<eapsDomain> | global]
```

Description

Displays summary EAPS counter information

Syntax Description

eapsDomain	Specifies the name of an EAPS domain. The switch displays counter information for only that domain.
global	Displays EAPS counter information when the events counted are not applicable to any specific EAPS domain.

Default

N/A.

Usage Guidelines

If you specify the name of an EAPS domain, the switch displays counter information related to only that domain. If you specify the `global` keyword, the switch displays EAPS counter information when the events counted are not applicable to any specific EAPS domain. The output displayed is for all configured EAPS domains, not just one specific EAPS domain.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults.

Clearing the Counters. The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, use one of the following commands:

- `clear counters`
- `clear eaps counters`

Understanding the Output. The following table describes the significant fields and values in the display output of the `show eaps counters <eapsDomain>` command:

Field	Description
Rx-Health	Indicates the EAPS domain received EAPS Health PDUs.
Rx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain received EAPS RingUp-FlushFdb PDUs to flush the FDB.
Rx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain received EAPS RingDown-FlushFdb PDUs to flush the FDB.
Rx-Link-Down	Indicates the EAPS domain received EAPS Link-Down PDUs and took down the link.
Rx-Flush-Fdb	Indicates the EAPS domain received EAPS Flush-Fdb PDUs and flushed the FDB.

Field	Description
Rx-Suspend-Prefwd-Timer	Indicates the EAPS domain received EAPS Suspend-Preforward-Timer PDUs. NOTE: Switches running ExtremeWare send this PDU during an MSM failover. Switches running ExtremeXOS 10.1 or later do not send or receive this PDU.
Rx-Query-Link-Status	Indicates the EAPS domain received EAPS Query-Link-Status PDUs.
Rx-Link-Up	Indicates the EAPS domain received EAPS Link-Up PDUs and brought the link back up.
Rx-Unknown	Indicates the EAPS domain dropped unknown EAPS PDUs.
Rx-Another-Master	Indicates the EAPS domain dropped EAPS PDUs because there is another Master switch in the same EAPS domain.
Rx-Unconfigured-Port	Indicates the EAPS domain dropped EAPS PDUs because the ingress port is not configured to be a ring port for the EAPS domain and the corresponding control VLAN.
Rx-Health-Pdu-Pri-Port	Indicates the EAPS domain dropped EAPS Health PDUs because the primary port received them instead of the secondary port. NOTE: The secondary port of the Master switch must receive EAPS Health PDUs, not the primary port.
Tx-Health	Indicates the EAPS domain sent EAPS Health PDUs.
Tx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain sent EAPS RingUp-FlushFdb PDUs to flush the FDB.
Tx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain sent EAPS RingDown-FlushFdb PDUs to flush the FDB.
Tx-Link-Down	Indicates the EAPS domain sent EAPS Link-Down PDUs because the link went down.
Tx-Flush-Fdb	Indicates the EAPS domain sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed.
Tx-Suspend-Prefwd-Timer	Indicates the EAPS domain sent EAPS Suspend-Preforward-Timer PDUs. NOTE: Switches running ExtremeWare send this PDU during an MSM failover. Switches running ExtremeXOS 10.1 or later do not send or receive this PDU. This counter should remain at 0.
Tx-Query-Link-Status	Indicates the EAPS domain sent EAPS Query-Link-Status PDUs.
Tx-Link-Up	Indicates the EAPS domain sent EAPS Link-Up PDUs and the link is up.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the EAPS domain. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the EAPS domain was unable to send because of an error.
Fw-Link-Down	Indicates the number of EAPS Link-Down PDUs received by the EAPS domain and forwarded in slow path.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the EAPS domain and forwarded in slow path.
FW-Query-Link-Status	Indicates the number of EAPS Query-Link-Status PDUs received by the EAPS domain and forwarded in slow path.
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Er	Indicates the number of EAPS PDUs the EAPS domain was unable to forward in slow path because of an error.

**NOTE**

Rx and Fw counters—If a PDU is received, processed, and consumed, only the Rx counter will increment. If a PDU is forwarded in slow path, both the Rx counter and Fw counter will increment.

The following table describes the significant fields and values in the display output of the `show eaps counters` global command:

Field	Description
Rx-Failed	Indicates an error occurred when receiving packets from the Layer 2 forwarding engine.
Rx-Invalid-Vlan-Intf	Indicates that the VLAN interface for the incoming VLAN cannot be found.
Rx-Undersize-Pkt	Indicates the length of the packet is less than the length of the header.
Rx-Invalid-8021Q-Tag	Indicates the VlanTypeLength field in the Ethernet header does not match the default Ethernet value for the 802.1Q tag.
Rx-Invalid-SNAP-Type	Indicates an invalid Subnetwork Access Protocol (SNAP) value in the Ethernet header.
Rx-Invalid-OUI	Indicates the Organizational Unique Identifier (OUI) value in the Ethernet header does not match 00:E0:2B.
Rx-EEP-Unsupported-Version	Indicates an unsupported Extreme Encapsulation Protocol (EEP) version. The EEP version should be 1.
Rx-EEP-Invalid-Length	Indicates the length of the EEP header is greater than the length of the packet.
Rx-EEP-Checksum-Invalid	Indicates the EEP checksum is invalid.
Rx-Domain-Invalid	Indicates the control VLAN's incoming PDU is not associated with an EAPS domain.
Rx-Lif-Invalid	Indicates that EAPS is unable to determine the logical interface (LIF) for the ingress port.
Rx-Lif-Down	Indicates the LIF for the ingress port is in the Down state.
Tx-Failed	Indicates an error occurred when sending packets to the Layer 2 forwarding engine.

Example

The following command displays the counters for a specific EAPS domain named eaps1:

```
show eaps counters eaps1
```

The following is sample output from this command:

```
Counters for EAPS domain: eaps1
Rx Stats
  Rx-Health                : 0
  Rx-Ringup-Flushfdb       : 0
  Rx-Ringdown-Flushfdb     : 0
  Rx-Link-Down             : 0
  Rx-Flush-Fdb             : 0
  Rx-Suspend-Prefwd-Timer  : 0
```

```

Rx-Query-Link-Status      : 0
Rx-Link-Up                : 0
Rx Dropped
Rx-Unknown                : 0
Rx-Another-Master         : 0
Rx-Unconfigured-Port      : 0
Rx-Health-Pdu-Pri-Port    : 0

Tx Stats
Tx-Health                 : 5011
Tx-Ringup-Flushfdb        : 0
Tx-Ringdown-Flushfdb      : 0
Tx-Link-Down              : 0
Tx-Flush-Fdb              : 0
Tx-Suspend-Prefwd-Timer   : 0
Tx-Query-Link-Status      : 3342
Tx-Link-Up                : 0
Tx Dropped
Tx-Unknown                : 0
Tx-Transmit-Err           : 0

Fw Stats
Fw-Link-Down              : 0
Fw-Flush-Fdb              : 0
Fw-Query-Link-Status      : 0
Fw Dropped
Fw-Unknown                : 0
Fw-Transmit-Err           : 0

```

The following command displays the global EAPS counters:

```
show eaps counters global
```

The following is sample output from this command:

```

Global counters for EAPS:
Rx-Failed                 : 0
Rx-Invalid-Vlan-Intf      : 0
Rx-Undersize-Pkt          : 0
Rx-Invalid-SNAP-Type      : 0
Rx-Invalid-OUI            : 0
Rx-EEP-Unsupported-Version : 0
Rx-EEP-Invalid-Length     : 0
Rx-EEP-Checksum-Invalid   : 0
Rx-Domain-Invalid         : 0
Rx-Failed                 : 0
Rx-Lif-Invalid            : 0
Rx-Lif-Down               : 0
Tx-Failed                 : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show eaps counters shared-port

```
show eaps counters shared-port [global | <port> {segment-port <segport>
{<eapsDomain>}}]
```

Description

Displays summary EAPS shared port counter information.

Syntax Description

global	Displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.
port	Identifies the port number of the specified common link port.
segport	Identifies the segment port. The segment port is the other ring port of an EAPS domain that is not the shared-port.
eapsDomain	Specifies the name of the EAPS domain. If no EAPS domain is specified, all counters for all EAPS domains on the specified segment port are displayed.

Default

N/A.

Usage Guidelines

If the switch is configured for EAPS shared ports, use this command to display an array of counters associated with the EAPS shared port functionality.

If you specify the `global` keyword, the switch displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.

If you specify a particular EAPS shared port, the switch displays counter information related to only that shared port.

If you specify a particular EAPS segment port, the switch displays counter information related to only that segment port for the specified EAPS domain.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults.

Clearing the Counters. The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, including the shared port counters, use one of the following commands:

- `clear counters`
- `clear eaps counters`

Understanding the Output. The following table describes the significant fields and values in the display output of the `show eaps counters shared-port` global command:

Field	Description
Rx-Invalid-Instance	Displays the number of dropped EAPS shared-port PDUs because there is not a valid EAPS shared port instance for the incoming port.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instances.
Fw-Invalid-Instance	Displays the number of EAPS shared-port PDUs that could not be forwarded in slow path because the shared port instances could not find a valid EAPS shared port instance for the outgoing port.

The following table describes the significant fields and values in the display output of the `show eaps counters shared-port <port> segment-port <segport> <eapsDomain>` command:

Field	Description
Rx-Seg-Health	Indicates the shared port instance received EAPS shared ports Segment-Health-Check PDUs.
Rx-Path-Detect	Indicates the shared port instance received EAPS shared ports Path-Detect PDUs.
Rx-Flush-Notify	Indicates the shared port instance received EAPS shared ports Flush-Notify PDUs and flushed the FDB. If this PDU reaches a port of the shared ports pair that initiated the PDU, the shared port instance may terminate the PDU. Otherwise, the shared port instance forwards the PDU.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instance.
Rx-Seg-Health-Dropped	Displays the number of EAPS shared ports Segment-Health-Check PDUs dropped by the shared port instance. This counter increments if the Segment-Health-Check PDU returns to the sending switch. If that occurs, the switch drops the Segment-Health-Check PDU.
Rx-Path-Detect-Dropped	Displays the number of EAPS shared ports Path-Detect PDUs dropped by the shared port instance. This counter increments in the following situations: • If the packet's Fwd-id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port. • If the packet's Link-Id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port.
Rx-Flush-Notify-Dropped	Displays the number of EAPS shared ports Flush-Notify-Dropped PDUs dropped by the shared port instance. This counter increments in the following situations: • If the Flush-Notify-Dropped PDU returns to the sending switch. • If the packet's Fwd-Id matches the EAPS shared port's Link-Id and the port is not in the blocking state.
Rx-Dropped-Invalid-Port	Displays the number of EAPS shared ports PDUs dropped by the shared port instance because it does not exist.
Tx-Seg-Health	Indicates the shared port instance sent EAPS shared ports Segment-Health-Check PDUs.
Tx-Path-Detect	Indicates the shared port instance sent EAPS shared ports Path-Detect PDUs. NOTE: This counter appears under Common Link Port Stats and should always be 0.

Field	Description
Tx-Flush-Notify	Indicates the shared port instance sent EAPS shared ports Flush-Notify PDUs to flush the FDB. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Flush-Fdb	Indicates the shared port instance sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the shared port instance. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to send because of an error.
Fw-Seg-Health	Indicates the number of EAPS shared ports Segment-Health-Check PDUs received by the shared port instance and forwarded in slow path.
Fw-Path-Detect	Indicates the number of EAPS shared ports Path-Detect PDUs received by the shared port instance and forwarded in slow path.
Fw-Flush-Notify	Indicates the number of EAPS Flush-Notify PDUs received by the shared port instance and forwarded in slow path to flush the FDB.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the shared port instance and forwarded in slow path.
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to forward in slow path because of an error.

Example

The following command displays global, high-level counter information for EAPS shared port:

```
show eaps counters shared-port global
```

The following is sample output from this command:

```
Global counters for EAPS Shared-Ports:
```

```
Rx Dropped
```

```
  Rx-Invalid-Instance : 0
```

```
  Rx-Unknown          : 0
```

```
Fw Dropped
```

```
  Fw-Invalid-Instance : 0
```

The following example assumes that port 17 is configured as an EAPS shared port. The following command displays counter information the specified EAPS shared port:

```
show eaps counters shared-port 17
```

The following is sample output from this command:

```
Counters for EAPS Shared-Port 17:
```

```
Common Link Port Stats
```

```

Rx Stats
  Rx-Seg-Health      : 0
  Rx-Path-Detect     : 0
  Rx-Flush-Notify    : 0
Rx Dropped
Rx-Seg-Health-Dropped : 0
  Rx-Path-Detect-Dropped : 0
  Rx-Flush-Notify-Dropped : 0
  Rx-Dropped-Invalid-Port : 0

Tx Stats
  Tx-Seg-Health      : 0
  Tx-Path-Detect     : 0
  Tx-Flush-Notify    : 0
  Tx-Flush-Fdb       : 0
Tx Dropped
  Tx-Unknown         : 0
  Tx-Transmit-Err    : 0

Fw Stats
  Fw-Seg-Health      : 0
  Fw-Path-Detect     : 0
  Fw-Flush-Notify    : 0
Fw Dropped
  Fw-Unknown         : 0
  Fw-Transmit-Err    : 0

```

The following example assumes that port 1:2 is configured as an EAPS shared port and port 1:1 is a segment port. The following command displays counter information the specified EAPS shared port, segment port, and EAPS domain:

```
show eaps counters shared-port 1:2 segment-port 1:1 eaps1
```

The following is sample output from this command:

```

Counters for EAPS Shared-Port 1:2, Segment Port: 1:1, EAPS Domain: eaps1
Rx Stats
  Rx-Seg-Health      : 0
  Rx-Path-Detect     : 0
  Rx-Flush-Notify    : 0
  Rx-Seg-Health-Dropped : 0
  Rx-Path-Detect-Dropped : 0
  Rx-Flush-Notify-Dropped : 0
  Rx-Dropped-Invalid-Port : 0
Tx Stats
  Tx-Seg-Health      : 2275
  Tx-Path-Detect     : 0
  Tx-Flush-Notify    : 0
  Tx-Flush-Fdb       : 0
  Tx-Transmit-Err    : 0
  Tx-Unknown         : 0
Fw Stats
  Fw-Seg-Health      : 0
  Fw-Path-Detect     : 0
  Fw-Flush-Notify    : 0
Fw-Transmit-Err      : 0
  Fw-Unknown         : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

show eaps shared-port

```
show eaps shared-port {<port>} {detail}
```

Description

Displays shared-port information for one or more EAPS domains.

Syntax Description

port	Specifies a shared-port.
detail	Specifies to display the status of all segments and VLANs.

Default

N/A.

Usage Guidelines

If you enter the `show eaps shared-port` command without the `detail` keyword, the command displays a summary of status information for all configured EAPS shared ports.

If you specify an EAPS shared-port, the command displays information about that specific port. Otherwise, the command displays information about all of the shared-ports configured on the switch.

In an EAPS shared port configuration, the segment ports are sorted in ascending order based on their port number, not the order you add an EAPS domain to EAPS shared ports. This is particularly useful when planning your EAPS configuration.



NOTE

The order you add EAPS domains to EAPS shared ports is relevant if the EAPS domains have matching ring ports and participate in spatial reuse. In this case, the newly added EAPS domain is added to the list and sorted on its ring port AFTER all other EAPS domains with the matching ring port.

The benefit of sorting ports in ascending order is evident if a common link fails. The port with the lowest port number among the segment ports in the “UP” state becomes the “Active Open” port. When configuring your EAPS domain, keep the port numbers in mind. For high bandwidth links, utilize lower port numbers, and for low bandwidth links, utilize higher port numbers. This way, if a common link fails, the high bandwidth link is still available.

You can use the `detail` keyword to display more detailed status information about the segments and VLANs associated with each shared port.

The following table describes the significant fields and values in the display output of the `show eaps shared-port {<port>} {detail}` commands:

Field	Description
Shared Port	Displays the port number of the shared port.
Mode	Indicates whether the switch on either end of the common link is a controller or partner. The mode is configured by the user.
Link ID	The link ID is the unique common link identifier configured by the user.
Up	Displays one of the following: • Yes—Indicates that the link ID and the mode are configured. • No—Indicates that the link ID or the mode is not configured.
State	Displays one of the following states: • Idle—Shared-port instance is not running. • Ready—The EAPS shared-port instance is running, the neighbor can be reached, and the common link is up. • Blocking—The EAPS shared-port instance is running, the neighbor cannot be reached, or the common link is down. • Preforwarding—The EAPS shared-port instance is in a blocking state, and the common link came up. To prevent a superloop, a temporary blocking state is created before going into Ready state.
Domain Count	Indicates the number of EAPS domains sharing the common link.
VLAN Count	Indicates the total number of VLANs that are protected under the EAPS domains sharing this common link.
Nbr	• Yes—Indicates that the EAPS instance on the other end of the common link is configured with matching link ID and opposite modes. For example, if one end of the common link is configured as a controller, the other end must be configured as a partner. • Err—Indicates that the EAPS instance on the other end of the common link is configured with a matching link ID, but the modes are configured the same. For example, both modes are configured as controller, or both modes are configured as partner. • No—The neighbor on the other end of the common link cannot be reached. Indicates one or more of the following: - The switch on the other end of the common link is not running. - The shared port has not been created. - The link IDs on each side of the common link do not match. - The common link, and any other segment, between the controller and partner are not fully connected.
RB State	• None—This EAPS shared-port is not the “root blocker.” • Active—This EAPS shared-port is the “root blocker” and is currently active. • Inactive—This EAPS shared-port is the “root blocker” but is currently inactive.
RB ID	The ID of the root blocker. If the value is none, there are not two or more common-link failures.
Active Open (available with the detail keyword)	• None—Indicates that there is no Active-Open port on the VLAN. • Port #—Indicates the port that is Active-Open and is in a forwarding state.

Field	Description
Segment Timer expiry action	- Segment down—Specifies that if the controller or partner switch detect a down segment, that segment stays down and a query is not sent through the ring. The switch marks the segment status as "Down." - Send alert—Specifies that if the controller or partner switch detect a down segment, that switch keeps the segment up and sends a warning message to the log (default). The switch sends a trap alert and sets the failed flag [F].
Segment Port (available with the <code>detail</code> keyword or by specifying a shared port)	The segment port is the other ring port of an EAPS domain that is not the shared-port.
Status (available with the <code>detail</code> keyword or by specifying a shared port)	- Up—There is connectivity to the neighboring EAPS shared-port via this port. - Down—There is a break in the path to the neighboring EAPS shared-port via this port. - Blocking-Up—The path is Up, but due to the "root blocker" being in the Active state, this port is blocked to prevent a loop. - Blocking-Down—The "root blocker" is in the Active State; however, the path is Down. Because the path is Down, there is no need to block the "root blocker" port to prevent a loop. [F]—The segment timer has expired but has not received an explicit link-down notification. The port remains in the Up state, with the timer expired flag set to True.
EAPS Domain (available with the <code>detail</code> keyword or by specifying a shared port)	The EAPS domain having the segment port as one of its ring ports.
Vlan-port count (available with the <code>detail</code> keyword or by specifying a shared port)	The total number of VLANs being protected under this segment port.
Adjacent Blocking Id (available with the <code>detail</code> keyword or by specifying a shared port)	- None—The neighbor on this port is not reporting a Controller in the Blocking state. <Link-Id>—The neighbor on this port is a controller in the Blocking state with a link ID of <Link-Id>.
Segment RB Id (available with the <code>detail</code> keyword or by specifying a shared port)	- None—The neighbor on this port is not aware of a "root blocker" in the network. <RB-Id>—The neighbor on this port has determined that there is a "root blocker" in the network with a link ID of <RB-Id>.
Vlan (available with the <code>detail</code> keyword or by specifying a shared port)	Displays a list of VLANs protected by the segment port.
Virtual-port Status (available with the <code>detail</code> keyword or by specifying a shared port)	This information appears for the Controller, when it is in either the Blocking or Preforwarding state. - Active-Open—This VLAN or port is in the Forwarding state and has connectivity to the neighboring EAPS shared port via this port. - Open—This VLAN or port is in the Forwarding state but does not have connectivity to the neighboring EAPS shared port via this port. - Blocked—This VLAN or port is in the Blocking state to prevent a loop in the network. - Down—This port's link is down. - Active—At this moment, this VLAN or port is not being handled by EAPS shared port. Rather, this VLAN or port is being handled by the regular EAPS protocol.

Example

The following command displays shared-port information for an EAPS shared port instance configured as a partner while it is in the ready state:

```
show eaps shared-port
```

```
EAPS shared-port count: 1
```

```
-----
Shared-port  Mode          Link      Domain Vlan      RB      RB
              Id    Up State   count   count Nbr State   Id
-----
8:1           Partner      1    Y Ready    2        1    Yes None   None
Segment Timer expiry action: Send alert
-----
```

The following command displays shared-port information for an EAPS shared port instance configured as a controller while it is in the ready state:

```
EAPS shared-port count: 1
```

```
-----
Shared-port  Mode          Link      Domain Vlan      RB      RB
              Id    Up State   count   count Nbr State   Id
-----
10:1         Controller  1    Y Ready    2        1    Yes None   None
Segment Timer expiry action: Send alert
-----
```

The following command displays information for an EAPS shared port instance configured as a controller while in the blocking state, with the detail option specified:

```
show eaps shared-port detail
```

```
EAPS shared-port count: 1
```

```
-----
Shared-port  Mode          Link      Domain Vlan      RB      RB
              Id    Up State   count   count Nbr State   Id
-----
4:1           Controller  10    Y Blocking 2        1    Yes Active  10
Segment Timer expiry action: Send alert
  Segment Port: 5:7, Status: Blocking-Up
    EAPS Domain:          d1
    Vlan-port count:      1
    Adjacent Blocking Id: None
    Segment RB Id:        None
    Vlan                  Virtual-port Status
      p_1                  Blocked

  Segment Port: 2:11,      Status: Down
    EAPS Domain:          d2
    Vlan-port count:      1
    Adjacent Blocking Id: 20
    Segment RB Id:        None
    Vlan                  Virtual-port Status
      p_1                  Open
-----
```

```
Vlan: p_1,   Vlan-port count: 2,   Active Open: None
Segment Port      Virtual-port Status
      5:7          Blocked
      2:11         Open
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

show vlan eaps

```
show vlan <vlan_name> eaps
```

Description

Displays the EAPS configuration (control, partner, or not added to an EAPS domain) of a specific VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Use this command to see if the specified VLAN is associated with an EAPS domain.

The output of this command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

If a VLAN is a partner VLAN for more than one EAPS domain, all of the EAPS domains that the VLAN is a partner of appears in the output.

Example

The following command displays the EAPS configuration for the control VLAN *orange* in EAPS domain *eaps1*:

```
show vlan orange eaps
```

The following is sample output from this command:

```
Vlan is Control in following EAPS domain:
    eaps1
```

The following command displays the EAPS configuration for the protected VLAN *purple* in EAPS domain *eaps1*:

```
show vlan purple eaps
```

The following is sample output from this command:

```
Vlan is Protected in following EAPS domain(s):
    eaps1
```

The following command displays information about the VLAN *default* not participating in EAPS:

```
show vlan default eaps
```

The following is sample output from this command:

```
Vlan has not been added to any EAPS domain
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure eaps shared-port link-id

```
unconfigure eaps shared-port <ports> link-id
```

Description

Unconfigures an EAPS link ID on a shared port on the switch.

Syntax Description

ports	Specifies the port number of the Common Link port.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the link ID on shared port 1:1.

```
unconfigure eaps shared-port 1:1 link-id
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

unconfigure eaps shared-port mode

```
unconfigure eaps shared-port <ports> mode
```

Description

Unconfigures the EAPS shared port mode.

Syntax Description

ports	Specifies the port number of the Common Link port.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the shared port mode on port 1:1:

```
unconfigure eaps shared-port 1:1 mode
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with a Core or an Advanced Core license.

unconfigure eaps port

```
unconfigure eaps <eapsDomain> [primary | secondary] port
```

Description

Sets the specified port's internal configuration state to INVALID.

Syntax Description

eapsDomain	Specifies the name of an EAPS domain.
primary	Specifies that the primary port should be unconfigured.
secondary	Specifies that the secondary port should be unconfigured.

Default

N/A.

Usage Guidelines

Unconfiguring an EAPS port sets its internal configuration state to INVALID, which causes the port to appear in the Idle state with a port status of Unknown when you use the `show eaps detail` command to display the status information about the port.

To prevent loops in the network, the switch displays by default a warning message and prompts you to unconfigure the specified EAPS primary or secondary ring port. When prompted, do one of the following:

- Enter `y` to unconfigure the specified port.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command on [page 1362](#).

Example

The following command unconfigures this node's EAPS primary ring port on the domain `eaps_1`:

```
unconfigure eaps eaps_1 primary port
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Unconfiguring the Primary port from the EAPS domain could cause a loop in
the network!
Are you sure you want to unconfigure the Primary EAPS Port? (y/n)
```

Enter `y` to continue and unconfigure the EAPS primary ring port. Enter `n` to cancel this action.

The switch displays a similar warning message if you unconfigure the secondary EAPS port.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

This chapter describes commands for:

- Creating, configuring, enabling, and disabling Spanning Tree Protocol (STP) on the switch
- Enabling and disabling Rapid Spanning Tree Protocol (RSTP) on the switch
- Enabling and disabling Multiple Spanning Tree Protocol (MSTP) on the switch
- Displaying and resetting STP settings on the switch

STP

The Spanning Tree Protocol (STP) is a bridge-based mechanism for providing fault tolerance on networks. STP is a part of the 802.1D bridge specification defined by the IEEE Computer Society. To explain STP in terms used by the 802.1D specification, the switch will be referred to as a bridge.

STP allows you to implement parallel paths for network traffic and ensure that redundant paths are:

- Disabled when the main paths are operational.
- Enabled if the main path fails.



NOTE

STP and Extreme Standby Router Protocol (ESRP) cannot be configured on the same Virtual LAN (VLAN) simultaneously.

RSTP

The Rapid Spanning Tree Protocol (RSTP) IEEE 802.1w provides an enhanced spanning tree algorithm that improves the convergence speed of bridged networks. RSTP takes advantage of point-to-point links in the network and actively confirms that a port can safely transition to the forwarding state without relying on any timer configurations. If a network topology change or failure occurs, RSTP rapidly recovers network connectivity by confirming the change locally before propagating that change to other devices across the network. For broadcast links, there is no difference in convergence time between STP and RSTP.

RSTP supersedes legacy STP protocols, supports the existing STP parameters and configurations, and allows for seamless interoperability with legacy STP.

MSTP

MSTP logically divides a Layer 2 network into regions. Each region has a unique identifier and contains multiple spanning tree instances (MSTIs). An MSTI is a spanning tree domain that operates within and is bounded by a region. MSTIs control the topology inside the regions. The Common and Internal Spanning Tree (CIST) is a single spanning tree domain that interconnects MSTP regions. The CIST is responsible for creating a loop-free topology by exchanging and propagating BPDUs across regions to form a Common Spanning Tree (CST).

MSTP uses RSTP as its converging algorithm and is interoperable with the legacy STP protocols: STP (802.1D) and RSTP (802.1w).

Spanning Tree Domains

The switch can be partitioned into multiple virtual bridges. Each virtual bridge can run an independent Spanning Tree instance. Each Spanning Tree instance is called a Spanning Tree Domain (STPD). Each STPD has its own root bridge and active path. After an STPD is created, one or more VLANs can be assigned to it.

A port can belong to multiple STPDs. In addition, a VLAN can span multiple STPDs.

The key points to remember when configuring VLANs and STP are:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

Member VLANs

When you add a VLAN to an STPD, that VLAN becomes a member of the STPD. The two types of member VLANs in an STPD are:

- Carrier
- Protected

Carrier VLAN. A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and if configured, the 802.1Q tag used to transport Extreme Multiple Instance Spanning Tree Protocol (EMISTP) or Per VLAN Spanning Tree (PVST+) encapsulated Bridge Protocol Data Units (BPDUs). Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.



NOTE

If you use EMISTP or PVST+, the StpdID must be identical to the VLANid of the carrier VLAN in that STPD.

If you have an 802.1D configuration, Extreme Networks recommends that you configure the StpdID to be identical to the VLANid of the carrier VLAN in that STPD.

If you configure MSTP, you do not need carrier VLANs for MSTP operation. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN. Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPD, but any particular port in the VLAN can belong to only *one* STPD. Also known as non-carrier VLANs.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the

CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

STPD Modes

An STPD has three modes of operation:

- 802.1D mode

Use this mode for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. When configured in this mode, all rapid configuration mechanisms are disabled.

- 802.1w mode

Use this mode for compatibility with Rapid Spanning Tree (RSTP). When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of this mode is available on point-to-point and edge ports only.

You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

- MSTP mode

Use this mode for compatibility with Multiple Spanning Tree (MSTP, 802.1s). MSTP is an extension of RSTP and offers the benefit of better scaling with fast convergence. When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of MSTP is available only on point-to-point links and when you configure the peer in MSTP or 802.1w mode. If you do not select point-to-point links and the peer is not configured in 802.1w mode, the STPD fails back to 802.1D mode.

You can create only one MSTP region on the switch, and all switches that participate in the region must have the same regional configurations. You enable or disable an MSTP on a per STPD basis only. You do not enable MSTP on a per port basis.

By default, the:

- STPD operates in 802.1D mode.
- Default device configuration contains a single STPD called *s0*.
- Default VLAN is a member of STPD *s0* with autobind enabled.

All STP parameters default to the IEEE 802.1D values, as appropriate.

Encapsulation Modes

You can configure ports within an STPD to accept and transmit specific BPDU encapsulations. This STP port encapsulation is separate from the STP mode of operation. For example, you can configure a port to accept the PVST+ BPDU encapsulation while running in 802.1D mode.

An STP port has three possible encapsulation modes:

- 802.1D mode

This mode is used for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- Extreme Multiple Instance Spanning Tree Protocol (EMISTP) mode

EMISTP mode is proprietary to Extreme Networks and is an extension of STP that allows a physical port to belong to multiple STPDs by assigning the port to multiple VLANs. EMISTP adds significant flexibility to STP network design. BPDUs are sent with an 802.1Q tag having an STPD instance Identifier (StpdID) in the VLANid field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- Per VLAN Spanning Tree (PVST+) mode

This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains to which it belongs.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STP Rules and Restrictions

This section summarizes the rules and restrictions for configuring STP as follows:

- The carrier VLAN must span all ports of the STPD. (This is not applicable to MSTP.)
- The StpdID must be the VLANid of the carrier VLAN; the carrier VLAN cannot be partitioned. (This is not applicable to MSTP.)
- A default VLAN cannot be partitioned. If a VLAN traverses multiple STPDs, the VLAN must be tagged.
- An STPD can carry, at most, one VLAN running in PVST+ mode, and its StpdID must be identical with that VLANid. In addition, the PVST+ VLAN cannot be partitioned.
- The default VLAN of a PVST+ port must be identical with the native VLAN on the PVST+ device connected to that port.
- If an STPD contains both PVST+ and non-PVST+ ports, that STPD must be enabled. If that STPD is disabled, the BPDUs are flooded in the format of the incoming STP port, which may be incompatible with those of the connected devices.
- The 802.1D ports must be untagged; and the EMISTP/PVST+ ports must be tagged in the carrier VLAN.
- An STPD with multiple VLANs must contain only VLANs that belong to the same virtual router instance.
- Automatically adding ports to an STPD (known as STP autobind) cannot be configured on a Netlogin VLAN.
- STP cannot be configured on the following ports:
 - A mirroring target port.
 - A software-controlled redundant port.
 - A Netlogin port.
- MSTP and 802.1D STPDs cannot share a physical port.
- Only one MSTP region can be configured on a switch.

- In an MSTP environment, A VLAN can belong to either a CIST or one of the MSTIs.
- A VLAN can belong to only one MSTP domain.
- MSTP is not interoperable with PVST+.
- The CIST can operate without any member VLANs.

clear counters stp

```
clear counters stp {[all | diagnostics | domains | ports]}
```

Description

Clears, resets all STP statistics and counters.

Syntax Description

all	Specifies all STP domain, port, and diagnostics counters.
diagnostics	Specifies STP diagnostics counters.
domains	Specifies STP domain counters.
ports	Specifies STP port counters.

Default

N/A.

Usage Guidelines

If you do not enter a parameter, the result is the same as specifying the `all` parameter: the counters for all domains, ports, and diagnostics are reset.

Enter one of the following parameters to reset the STP counters on the switch:

- `all`—Specifies the counters for all STPDs and ports, and clears all STP counters
- `diagnostics`—Clears the internal diagnostic counters
- `domains`—Clears the domain level counters
- `ports`—Clears the counters for all ports and leaves the domain level counters

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period that you are monitoring.

Example

The following command clears all of the STP domain, port, and diagnostic counters:

```
clear counters stp
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure mstp format

```
configure mstp format <format_identifier>
```

Description

Configures the number used to identify the MSTP BPDUs sent in the MSTP region.

Syntax Description

format_identifier	Specifies a number that MSTP uses to identify all BPDUs sent in the MSTP region. The default is 0. The range is 0 to 255.

Default

The default value used to identify the MSTP BPDU is 0.

Usage Guidelines

For a switch to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches contained in a region transmit and receive BPDUs that contain information relevant to only that MSTP region. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, Extreme Networks recommends that you disable all active STPDs in the region before modifying the value used to identify MSTP BPDUs on all participating switches.

Example

The following command configures the number 2 to identify the MSTP BPDUs sent within an MSTP region:

```
configure mstp format 2
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mstp region

```
configure mstp region <regionName>
```

Description

Configures the name of an MSTP region on the switch.

Syntax Description

regionName	Specifies a user-defined name for the MSTP region. May be up to 32 characters.

Default

By default, the switch uses the MAC address of the switch to generate an MSTP region.

Before you configure the MSTP region, it also has the following additional defaults:

- MSTP format Identifier—0
- MSTP Revision Level—3

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, mstp. Names must start with an alphabetical character, for example, a, Z.

By default, the switch uses the unique MAC address of the switch to generate an MSTP region. Since each MAC address is unique, every switch is in its own region by default.

For multiple switches to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches inside a region exchange BPDUs that contain information for MSTIs. The switches connected outside of the region exchange CIST information. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, Extreme Networks recommends that you disable all active STPDs in the region before renaming the region on all of the participating switches.

Viewing MSTP Information. To view the MSTP configuration on the switch, use the `show stpd` command. Output from this command contains global MSTP settings, including the name of the MSTP

region, the number or tag that identifies all of the BPDUs sent in the MSTP region, and the reserved MSTP revision level. If configured, the output also displays the name of the Common and Internal Spanning Tree (CIST), and the number of Multiple Spanning Tree Instances (MSTIs).

Example

The following command creates an MSTP region named purple:

```
configure mstp region purple
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mstp revision

```
configure mstp revision <revision>
```

Description

Configures the revision number of the MSTP region.

Syntax Description

revision	This parameter is reserved for future use.

Default

The default value of the revision level is 3.

Usage Guidelines

Although this command is displayed in the CLI, it is reserved for future use. Please do not use this command.

If you accidentally configure this command, remember that each switch in the region must have the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command returns the MSTP revision number to 3, the default revision number:

```
configure mstp revision 3
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd add vlan

```
configure stpd <stpd_name> add vlan <vlan_name> ports [all | <port_list>]
{[dot1d | emistp | pvst-plus]}
```

Description

Adds all ports or a list of ports within a VLAN to a specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies a VLAN name.
all	Specifies all of the ports in the VLAN to be included in the STPD.
port_list	Specifies the port or ports to be included in the STPD.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1D.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

To create an STP domain, use the `create stpd` command. To create a VLAN, use the `create vlan` command.

In an EMISTP or PVST+ environment, this command adds a list of ports within a VLAN to a specified STPD provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the `dot1d` encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports
- Software-controlled redundant ports
- Netlogin ports

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- Another 802.1D mode STP port to a physical port that already contains an 802.1D mode STP port (only *one* 802.1D encapsulation STP port can be configured on a particular STP port)
- A carrier VLAN port to a different STP domain than the carrier VLAN belongs
- A VLAN/port for which the carrier VLAN does not yet belong



NOTE

This restriction is only enforced in an active STP domain and when you enable STP to make sure you have a legal STP configuration.

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named *default* with a tag value of 1 and STPD *s0*. The switch associates VLAN *default* to STPD *s0*. All ports that belong to this VLAN and STPD are in 802.1D encapsulation mode with autobind enabled. If you disable autobind on the VLAN *default*, that configuration is saved across a reboot.

Naming Conventions. If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

STP Encapsulations Modes. You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an StpdID in the VLANid field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STPD Identifier. An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An StpdID must be identical to the VLANid of the carrier VLAN in that STPD and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD

that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Automatically Inheriting Ports—MSTP Only. In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Example

Create a VLAN named *marketing* and an STPD named *STPD1* as follows:

```
create vlan marketing
create stpd stpd1
```

The following command adds the VLAN named *marketing* to the STPD *STPD1*, and includes all the ports of the VLAN in *STPD1*:

```
configure stpd stpd1 add vlan marketing ports all
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd default-encapsulation

```
configure stpd <stpd_name> default-encapsulation [dot1d | emistp | pvst-plus]
```

Description

Configures the default encapsulation mode for all ports added to the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named *default* with a tag value of 1 and STPD *s0*. The switch associates VLAN *default* to STPD *s0*. All ports that belong to this VLAN and STPD are in 802.1d encapsulation mode with autobind enabled. If you disable autobind on the VLAN *default*, that configuration is saved across a reboot.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Naming Conventions. If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

STP Encapsulation Modes. You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an StpdID in the VLANid field. This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- **pvst-plus**—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.



NOTE

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

STPD Identifier. An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An StpdID must be identical to the VLANid of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Example

The following command specifies that all ports subsequently added to the STPD *STPD1* be in PVST+ encapsulation mode unless otherwise specified or manually changed:

```
configure stpd stpd1 default-encapsulation pvst-plus
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd delete vlan

```
configure stpd <stpd_name> delete vlan <vlan_name> ports [all |  
<port_list>]
```

Description

Deletes one or more ports in the specified VLAN from an STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies a VLAN name.
all	Specifies all of the ports in the VLAN to be removed from the STPD.
port_list	Specifies the port or ports to be removed from the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

In EMISTP and PVST+ environments, if the specified VLAN is the carrier VLAN, all protected VLANs on the same set of ports are also removed from the STPD.

You also use this command to remove autobind ports from a VLAN. ExtremeXOS records the deleted ports so that the ports are not automatically added to the STPD after a system restart.

When a port is deleted on the MSTI, it is automatically deleted on the CIST as well.

Example

The following command removes all ports of a VLAN named *Marketing* from the STPD *STPD1*:

```
configure stpd stpd1 delete vlan marketing ports all
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd forwarddelay

```
configure stpd <stpd_name> forwarddelay <seconds>
```

Description

Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the Root Bridge.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
seconds	Specifies the forward delay time in seconds. The default is 15 seconds, and the range is 4 to 30 seconds.

Default

The default forward delay time is 15 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<seconds>` parameter is 4 through 30 seconds.

Example

The following command sets the forward delay from *STPD1* to 20 seconds:

```
configure stpd stpd1 forwarddelay 20
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd hellotime

```
configure stpd <stpd_name> hellotime <seconds>
```

Description

Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the Root Bridge.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
seconds	Specifies the hello time in seconds. The default is 2 seconds, and the range is 1 to 10 seconds.

Default

The default hello time is 2 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

In an MSTP environment, configure the hello timer only on the CIST, not on the MSTIs.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<seconds>` parameter is 1 through 10 seconds.

Example

The following command sets the time delay from *STPD1* to 10 seconds:

```
configure stpd stpd1 hellotime 10
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd maxage

```
configure stpd <stpd_name> maxage <seconds>
```

Description

Specifies the maximum age of a BPDU in the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
seconds	Specifies the maxage time in seconds. The default is 20 seconds, and the range is 6 to 40 seconds.

Default

The default maximum age of a BPDU is 20 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

In an MSTP environment, configure the maximum age of a BPDU only on the CIST, not on the MSTIs.

The range for the `<seconds>` parameter is 6 through 40 seconds.

Note that the time must be greater than, or equal to $2 * (\text{Hello Time} + 1)$ and less than, or equal to $2 * (\text{Forward Delay} - 1)$.

Example

The following command sets the maximum age of *STPD1* to 30 seconds:

```
configure stpd stpd1 maxage 30
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd max-hop-count

```
configure stpd <stpd_name> max-hop-count <hopcount>
```

Description

Specifies the maximum hop count of a BPDU until the BPDU is discarded in the specified MSTP STP domain.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
hopcount	Specifies the number of hops required to age out information and notify changes in the topology. The default is 20 hops, and the range is 6 to 40 hops.

Default

The default hop count of a BPDU is 20 hops.

Usage Guidelines

This command is applicable only in an MSTP environment.

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<hopcount>` parameter is 6 through 40 hops.

In an MSTP environment, the hop count has the same purpose as the maxage timer for 802.1D and 802.1w environments.

The main responsibility of the CIST is to exchange or propagate BPDUs across regions. The switch assigns the CIST an instance ID of 0, which allows the CIST to send BPDUs for itself in addition to all of the MSTIs within an MSTP region. Inside a region, the BPDUs contain CIST records and piggybacked M-records. The CIST records contain information about the CIST, and the M-records contain information about the MSTIs. Boundary ports only exchange CIST record BPDUs.

On boundary ports, only CIST record BPDUs are exchanged. In addition, if the other end is an 802.1D or 802.1w bridge, the maxage timer is used for interoperability between the protocols.

Example

The following command sets the hop of the MSTP STPD, STPD2, to 30 hops:

```
configure stpd stpd2 max-hop-count 30
```


History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd mode

```
configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti
<instance>]]
```

Description

Configures the operational mode for the specified STP domain.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STPD mode of operation to be 802.1D.
dot1w	Specifies the STPD mode of operation to be 802.1w, and rapid configuration is enabled.
mstp	Specifies the STPD mode of operation to be 802.1s, and rapid configuration is enabled.
cist	Configures the specified STPD as the common instance spanning tree for the MSTP region.
msti	Configures the specified STPD as a multiple spanning tree instance for the MSTP region.
instance	Specifies the Id of the multiple spanning tree instance. The range is 1 to 4,094.

Default

The STPD operates in 802.1d mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If you configure the STP domain in 802.1D mode, the rapid reconfiguration mechanism is disabled.

If you configure the STP domain in 802.1w mode, the rapid reconfiguration mechanism is enabled. You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

If you configure the STP domain in MSTP mode, the rapid reconfiguration mechanism is enabled. You enable or disable MSTP on a per STPD basis only. You do not enable MSTP on a per port basis. MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You must first configure a Common and Internal Spanning Tree (CIST) before configuring any multiple spanning tree instances (MSTIs) in the region. You cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command configures STPD *s1* to enable the rapid reconfiguration mechanism and operate in 802.1w mode:

```
configure stpd s1 mode dot1w
```

The following command configures STPD *s2* to operate as an MSTI in an MSTP domain:

```
configure stpd s2 mode mstp msti 3
```

History

This command was available in ExtremeXOS 10.1.

The `mstp` parameter was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd ports cost

```
configure stpd <stpd_name> ports cost [auto | <cost>] <port_list>
```

Description

Specifies the path cost of the port in the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
auto	Specifies the switch to remove any user-defined port cost value(s) and use the appropriate default port cost value(s).
cost	Specifies a numerical port cost value. The range is 1 through 200,000,000.
port_list	Specifies one or more ports or slots and ports.

Default

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 2,000,000.
- 100 Mbps port—the default cost is 200,000.
- 1000 Mbps port—the default cost is 20,000.
- 10000 Mbps ports—the default cost is 2,000.

The default port cost for trunked ports is dynamically calculated based on the available bandwidth.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The 802.1D-2004 standard modified the default port path cost value to allow for higher link speeds. If you have a network with both 802.1D-2004 and 802.1D-1998 compliant bridges, a higher link speed can create a situation whereby an 802.1D-1998 compliant bridge could become the most favorable transit path and possibly cause the traffic to span more bridges. To prevent this situation, configure the port path cost to make links with the same speed use the same path host value. For example, if you have 100 Mbps links on all bridges, configure the port path cost for the 802.1D-2004 compliant bridges to 19 instead of using the default 200,000.



NOTE

You cannot configure the port path cost on 802.1D-1998 compliant bridges to 200,000 because the path cost range setting is 1 to 65,535.

The range for the `cost` parameter is 1 through 200,000,000. If you configure the port cost, a setting of 1 indicates the highest priority.

If you configured a port cost value and specify the `auto` option, the switch removes the user-defined port cost value and returns to the default, automatically assigned, port cost value.

The `auto` port cost of a trunk port is calculated based on number member ports in the trunk port. Link up and down of the member port does not affect the trunk port cost, thus it does not trigger topology change. Only adding or removing a member port to/from the trunk port causes `auto` trunk port cost to change. Also, by so configuring a static trunk port cost, the value is frozen regardless of the number of member ports in the trunk port.

ExtremeXOS 11.5 and Earlier. If you have switches running ExtremeXOS 11.5 and earlier, the default costs are different than switches running ExtremeXOS 11.6 and later.

The range for the `cost` parameter is 1 through 65,535.

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 100.
- 100 Mbps port—the default cost is 19.
- 1000 Mbps port—the default cost is 4.
- 10000 Mbps ports—the default cost is 2.

Example

The following command configures a cost of 100 to slot 2, ports 1 through 5 in STPD `s0`:

```
configure stpd s0 ports cost 100 2:1-2:5
```

History

This command was available in ExtremeXOS 10.1.

The `auto` option was added in ExtremeXOS 11.0.

The default costs were updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd ports edge-safeguard disable

```
configure stpd <stpd_name> ports edge-safeguard disable <port_list>
```

Description

Disables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more edge ports.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called “edge safeguard.” An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs.

If you disable this feature, the edge port enters the forwarding state but no longer transmits BPDUs unless a BPDU is received by that edge port. This is the default behavior.

To view the status of the edge safeguard feature use the `show stpd <stpd_name> ports {[detail | <port_list> {detail}}]` command. You can also use the `show stpd {<stpd_name> | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.

To enable or re-enabled edge safeguard, use one of the following commands:

- `configure stpd <stpd_name> ports edge-safeguard enable <port_list>`
- `configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> | edge <port_list> {edge-safeguard [enable | disable]}]`

Example

The following command disables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard disable 4
```

The following command disables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD S1 on a modular switch:

```
configure stpd s1 ports edge-safeguard disable 2:3
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd ports edge-safeguard enable

```
configure stpd <stpd_name> ports edge-safeguard enable <port_list>
```

Description

Enables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more edge ports.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called “edge safeguard.” You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

To view the status of the edge safeguard feature use the `show stpd <stpd_name> ports {[detail | <port_list> {detail}}]` command. You can also use the `show stpd {<stpd_name> | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.

To disable edge safeguard, use one of the following commands:

- `configure stpd <stpd_name> ports edge-safeguard disable <port_list>`
- `configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> | edge <port_list> {edge-safeguard [enable | disable]}]`

Example

The following command enables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard enable 4
```


The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD S1 on a modular switch:

```
configure stpd s1 ports edge-safeguard enable 2:3
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd ports link-type

```
configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> | edge <port_list> {edge-safeguard [enable | disable]}]
```

Description

Configures the ports in the specified STPD as auto, broadcast, edge, or point-to-point link types.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
auto	Specifies the switch to automatically determine the port link type. An auto link behaves like a point-to-point link if the link is in full-duplex mode or if link aggregation is enabled on the port. Used for 802.1w configurations.
broadcast	Specifies a port attached to a LAN segment with more than two bridges. Used for 802.1D configurations. A port with broadcast link type cannot participate in rapid reconfiguration using RSTP or MSTP. By default, all STP.1D ports are broadcast links.
point-to-point	Specifies a port attached to a LAN segment with only two bridges. A port with point-to-point link type can participate in rapid reconfiguration. Used for 802.1w and MSTP configurations. By default, all 802.1w and MSTP ports are point-to-point link types.
port_list	Specifies one or more ports or slots and ports.
edge	Specifies a port that does not have a bridge attached. An edge port is placed and held in the STP forwarding state unless a BPDU is received by the port. Used for 802.1w and MSTP configurations.
edge-safeguard	Specifies that the edge port be configured with edge safeguard, a loop prevention and detection mechanism. Used for 802.1w and MSTP configurations
enable	Specifies that edge safeguard be enabled on the edge port(s).
disable	Specifies that edge safeguard be disabled on the edge port(s).

Default

STP.1D ports are broadcast link types
802.1w and MSTP ports are point-to-point link types

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

The default, broadcast links, supports legacy STP (802.1D) configurations. If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

RSTP rapidly moves the designated ports of a point-to-point link type into the forwarding state. This behavior is supported by RSTP and MSTP only.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Auto Link Type. An auto link behaves like a point-to-point link if the link is in full duplex mode or if link aggregation is enabled on the port; otherwise, an auto link behaves like a broadcast link. If a non-STP switch exists between several switches operating in 802.1w mode with auto links, the non-STP switch may negotiate full-duplex even though the broadcast domain extends over several STP devices.

Edge Link Type. RSTP does not send any BPDUs from an edge port nor does it generate topology change events when an edge port changes its state.

If you configure a port to be an edge port, the port immediately enters the forwarding state. Edge ports remain in the forwarding state unless the port receives a BPDU. In that case, edge ports enter the blocking state. The edge port remains in the blocking state until it stops receiving BPDUs and the message age timer expires.

Edge Safeguard. Loop prevention and detection on an edge port configured for RSTP or MSTP is called “edge safeguard.” You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

To enable configure a port as an edge port and enable edge safeguard on that port, use the `configure stpd <stpd_name> ports link-type edge <port_list> edge-safeguard enable` command and specify `enable`.

To disable edge safeguard on the edge port, use the `configure stpd <stpd_name> ports link-type edge <port_list> edge-safeguard disable` command and specify `disable`.

Two other commands are also available to enable and disable edge safeguard:

```
configure stpd ports edge-safeguard enable
configure stpd ports edge-safeguard disable
```

Example

The following command configures slot 2, ports 1 through 4 to be point-to-point links in STPD s1:

```
configure stpd s1 ports link-type point-to-point 2:1-2:4
```

The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 configured for RSTP:

```
configure stpd s1 ports link-type edge 2:3 edge-safeguard enable
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd ports mode

```
configure stpd <stpd_name> ports mode [dot1d | emistp | pvst-plus]
<port_list>
```

Description

Configures the encapsulation mode for the specified port list.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.
port_list	Specifies one or more ports or slots and ports.

Default

Ports in the default STPD (s0) are `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.
This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.
- `emistp`—This mode sends BPDUs with an 802.1Q tag having an StpdID in the VLANid field.
This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.
- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.
This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

Example

The following command configures STPD *s1* with PVST+ packet formatting for slot 2, port 1:

```
configure stpd s1 ports mode pvst-plus 2:1
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd ports port-priority

```
configure stpd <stpd_name> ports port-priority <priority> <port_list>
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
priority	Specifies a numerical port priority value. The range is 0 through 240 and is subject to the multiple of 16 restriction.
port_list	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing [configure stpd ports priority](#) command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry. For example, if the switch reads the `configure stpd ports priority 16` command from an ExtremeXOS 11.5 or earlier configuration, (which is equivalent to the command `configure stpd ports priority 8` entered through CLI), the switch saves the value in the new ExtremeXOS 11.6 configuration as `configure stpd ports port-priority 128`.

A setting of 0 indicates the highest priority.

The range for the `priority` parameter is 0 through 240 and is subject to the multiple of 16 restriction.

Example

The following command assigns a priority of 32 to slot 2, ports 1 through 5 in STPD `s0`:

```
configure stpd s0 ports port-priority 32 2:1-2:5
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd ports priority

```
configure stpd <stpd_name> ports priority <priority> <port_list>
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
priority	Specifies a numerical port priority value. The range is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.
port_list	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing [configure stpd ports priority](#) command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry.

A setting of 0 indicates the highest priority.

The range for the `priority` parameter is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.

ExtremeXOS 11.6 introduces support for a new ports priority command: [configure stpd ports port-priority](#). When you save the port priority value in an ExtremeXOS 11.6 configuration, the switch saves it as the new command [configure stpd ports port-priority](#) with the corresponding change in priority values. The priority range of this command is 0 through 240 and is subject to the multiple of 16 restriction. For more information see [“configure stpd ports port-priority” on page 1450](#).

ExtremeXOS 11.5 and Earlier. If you have switches running ExtremeXOS 11.5 and earlier, the default value for the priority range are different than switches running ExtremeXOS 11.6.

The range for the `priority` parameter is 0 through 31.

The default is 16.

Example

The following command assigns a priority of 1 to slot 2, ports 1 through 5 in STPD *s0*:

```
configure stpd s0 ports priority 1 2:1-2:5
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

configure stpd priority

```
configure stpd <stpd_name> priority <priority>
```

Description

Specifies the bridge priority of the STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
priority	Specifies the bridge priority of the STPD. The range is 0 through 61,440 and is subject to the multiple of 4,096 restriction.

Default

The default priority is 32,768.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the bridge priority of the STPD, you can make it more or less likely to become the root bridge.

The range for the `<priority>` parameter is 0 through 61,440 and is subject to the multiple of 4,096 restriction. A setting of 0 indicates the highest priority.

If you have an ExtremeXOS 11.5 or earlier configuration that contains an STP or RSTP bridge priority that is not a multiple of 4,096, the switch rejects the entry and the bridge priority returns to the default value. The MSTP implementation already uses multiples of 4,096 to determine the bridge priority.

For example, to lower the numerical value of the priority (which gives the priority a higher precedence), you subtract 4,096 from the default priority: $32,768 - 4,096 = 28,672$. If you modify the priority by a value other than 4,096, the switch rejects the entry.

ExtremeXOS 11.5 and Earlier. If you have switches running ExtremeXOS 11.5 and earlier, the priority range is different than switches running ExtremeXOS 11.6 and later.

The range for the `priority` parameter is 0 through 65,535. A setting of 0 indicates the highest priority.

Example

The following command sets the bridge priority of *STPD1* to 16,384:

```
configure stpd stpd1 priority 16384
```

History

This command was available in ExtremeXOS 10.1.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd tag

```
configure stpd <stpd_name> tag <stpd_tag>
```

Description

Assigns an StpdID to an STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
stpd_tag	Specifies the VLANid of the carrier VLAN that is owned by the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An StpdID must be identical to the VLANid of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD. Unless all ports are running in 802.1D mode, an STPD with ports running in either EMISTP mode or PVST+ mode must be configured with an StpdID.

You must create and configure the VLAN, along with the tag, before you can configure the STPD tag. To create a VLAN, use the `create vlan` command. To configure the VLAN, use the `configure vlan` commands.

MSTP Only. MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the CIST. The switch assigns this ID automatically when you configure the CIST STPD. To configure the CIST STPD, use the `configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti <instance>]]` command.

An MSTI identifier (MSTI ID) identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. Each STPD that participates in a particular MSTP region must have the same MSTI ID. To configure the MSTI ID, use the `configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti <instance>]]` command.

Example

The following command assigns an StpdID to the `purple_st` STPD:

```
configure stpd purple_st tag 200
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan add ports stpd

```
configure vlan <vlan_name> add ports [all | <port_list>] {tagged |
untagged} stpd <stpd_name> {[dot1d | emistp | pvst-plus]}
```

Description

Adds one or more ports in a VLAN to a specified STPD.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all of the ports to be included in the STPD.
port_list	Specifies the port or ports to be included in the STPD.
tagged	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

To create a VLAN, use the `create vlan` command. To create an STP domain, use the `create stpd` command.

In an EMISTP or PVST+ environment, this command adds a list of ports to a VLAN and a specified STPD at the same time provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the `dot1d` encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports
- Software-controlled redundant ports
- Netlogin ports

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- Another 802.1D mode STP port to a physical port that already contains an 802.1D mode STP port (only one 802.1D encapsulation STP port can be configured on a particular STP port)
- A carrier VLAN port to a different STP domain than the carrier VLAN belongs
- A VLAN/port for which the carrier VLAN does not yet belong



NOTE

This restriction is only enforced in an active STP domain and when you enable STP to ensure you have a legal STP configuration.

Naming Conventions. If your VLAN has the same name as another component, for example an STPD, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your VLAN has a name unique only to that VLAN, the keywords `vlan` and `stpd` are optional.

STP Encapsulation Modes. You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an `StpdID` in the `VLANid` field. This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.
- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical ports belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Automatically Inheriting Ports—MSTP Only. In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Example

The following command adds slot 1, port 2 and slot 2, port 3, members of a VLAN named *Marketing*, to the STPD named *STPD1*, and specifies that they be in *EMISTP* mode:

```
configure vlan marketing add ports 1:2, 2:3 tagged stpd stpd1 mistp
```

History

This command was available in ExtremeXOS 10.1.

The nobroadcast keyword was removed in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

create stpd

```
create stpd <stpd_name>
```

Description

Creates a user-defined STPD.

Syntax Description

stpd_name	Specifies a user-defined STPD name to be created. May be up to 32 characters in length.

Default

The default device configuration contains a single STPD called *s0*.

When an STPD is created, the STPD has the following default parameters:

- State—disabled
- StpdID—none
- Assigned VLANs—none
- Bridge priority—32,768
- Maximum BPDU age—20 seconds
- Hello time—2 seconds
- Forward delay—15 seconds
- Operational mode—802.1D
- Rapid Root Failover—disabled
- Default Binding Mode (encapsulation mode)—Ports in the default STPD (s0) are in 802.1d mode. Ports in user-created STPDs are in *emistp* mode.
- Maximum hop count (when configured for MSTP)—20 hops

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (`_`) but cannot be any reserved keywords, for example, *stp* or *stpd*. Names must start with an alphabetical character, for example, *a*, *Z*.

Each STPD name must be unique and cannot duplicate any other named STPDs on the switch. If you are uncertain about the STPD names on the switch, use the `show stpd` command to view the STPD names.

You can, however, re-use names across multiple categories of switch configuration. For example, you can use the name *Test* for an STPD and a VLAN. If you use the same name, Extreme Networks recommends that you specify the appropriate keyword when configuring the STPD. If you do not specify the appropriate keyword, the switch displays a message similar to the following:

```
%% Ambiguous command:  "configure Test"
```

To view the names of the STPDs on the switch, enter `configure` and press [Tab]. Scroll to the end of the output to view the names.

Each STPD has its own Root Bridge and active path. After the STPD is created, one or more VLANs can be assigned to it.

Example

The following example creates an STPD named *purple_st*:

```
create stpd purple_st
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete stpd

```
delete stpd <stpd_name>
```

Description

Removes a user-defined STPD from the switch.

Syntax Description

stpd_name	Specifies a user-defined STPD name on the switch.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command:  "delete Test"
```

In this example, to delete the STPD *Test*, enter `delete stpd Test`.

If you created an STPD with a name unique only to that STPD, the keyword `stpd` is optional.

The default STPD, *s0*, cannot be deleted.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command deletes an STPD named *purple_st*:

```
delete stpd purple_st
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd

```
disable stpd {<stpd_name>}
```

Description

Disables the STP protocol on a particular STPD or for all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

After you have created the STPD with a unique name, the keyword *stpd* is optional.

If you want to disable the STP protocol for all STPDs, do not specify an STPD name.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command disables an STPD named *purple_st*:

```
disable stpd purple_st
```

The following command disables the STP protocol for all STPDs on the switch:

```
disable stpd
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd auto-bind

```
disable stpd <stpd_name> auto-bind vlan <vlan_name>
```

Description

Disables the ability to automatically add ports to an STPD when they are added to a member VLAN.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies the name of a member VLAN with autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

Usage Guidelines



NOTE

Ports already in the STPD remain in that domain (as if they were added manually).

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

Ports added to the STPD automatically when autobind is enabled are not removed when autobind is disabled. The ports are present after a switch reboot.

To view STP configuration status of the ports in a VLAN, use the following command:

```
show vlan <vlan_name> stpd
```

Example

The following example disables autobind on an STPD named `s8`:

```
disable stpd s8 auto-bind v5
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd ports

```
disable stpd <stpd_name> ports [all | <port_list>]
```

Description

Disables STP on one or more ports for a given STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
all	Specifies all ports for a given STPD.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create the STPD with a unique name, the keyword `stpd` is optional.

Disabling STP on one or more ports puts those ports in the *forwarding* state; all BPDUs received on those ports are disregarded and dropped.

Use the `all` keyword to specify that all ports of a given STPD are disabled.

Use the `port_list` parameter to specify a list of ports of a given STPD are disabled.

If you do not use the default STPD, you must create one or more STPDs and configure and enable the STPD before you can use the `disable stpd ports` command.

Example

The following command disables slot 2, port 4 on an STPD named *Backbone_st*:

```
disable stpd backbone_st ports 2:4
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd rapid-root-failover

```
disable stpd <stpd_name> rapid-root-failover
```

Description

Disables rapid root failover for STP recovery times.

Syntax Description

stpd_name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

After you have created the STPD with a unique name, the keyword `stpd` is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command disables rapid root fail over on STPD *Backbone_st*:

```
disable stpd backbone_st rapid-root-failover
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd

```
enable stpd {<stpd_name>}
```

Description

Enables the STP protocol for one or all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

If you want to enable the STP protocol for all STPDs, do not specify an STPD name.

Example

The following command enables an STPD named *Backbone_st*:

```
enable stpd backbone_st
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd auto-bind

```
enable stpd <stpd_name> auto-bind vlan <vlan_name>
```

Description

Automatically adds ports to an STPD when ports are added to a member VLAN.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies the name of the VLAN to have autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

If you enable autobind and add ports to a member VLAN, those ports are automatically added to the STPD.

Usage Guidelines

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

You cannot configure the autobind feature on a Netlogin VLAN.

In an EMISTP or PVST+ environment, when you issue this command, any port or list of ports that you add to the carrier VLAN are automatically added to the STPD with autobind enabled. In addition, any port or list of ports that you remove from a carrier VLAN are automatically removed from the STPD. This allows the STPD to increase or decrease its span as you add ports to or remove ports from a carrier VLAN.

For MSTP, when you issue this command, any port or list of ports that gets automatically added to an MSTI are automatically inherited by the CIST. In addition, any port or list of ports that you remove from an MSTI protected VLAN are automatically removed from the CIST. For more information see the section. For more information, see [“Automatically Inheriting Ports—MSTP Only” on page 1470](#).

Carrier VLAN. A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and the 802.1Q tag used to transport STP BPDUs in the encapsulation mode is EMISTP or PVST+. Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.



NOTE

The StpdID must be identical to the VLANid of the carrier VLAN in that STPD.

If you configure MSTP, you do not need a carrier VLAN. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN. Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPDs, but any particular port in the VLAN can belong to only one STPD.

Enabling autobind on a protected VLAN does not expand the boundary of the STPD. However, the VLAN and port combinations are added to or removed from the STPD subject to the boundaries of the carrier VLAN.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

Automatically Inheriting Ports—MSTP Only. In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Displaying STP Information. To view STP configuration status of the ports on a VLAN, use the following command:

```
show vlan <vlan_name> stpd
```

Example

The examples in this section assume that you have already removed the ports from the Default VLAN.

To automatically add ports to an STPD running 802.1D, EMISTP, or PVST+ and to expand the boundary of the STPD, you must complete the following tasks:

- Create the carrier VLAN.
- Assign a VLANid to the carrier VLAN.
- Add ports to the carrier VLAN.
- Create an STPD (or use the default, S0).
- Enable autobind on the STPDs carrier VLAN.
- Configure the STPD tag (the StpdID must be identical to the VLANid of the carrier VLAN in the STP domain).
- Enable STP.

The following example enables autobind on an STPD named *s8* after creating a carrier VLAN named *v5*:

```
create vlan v5
configure vlan v5 tag 100
configure vlan v5 add ports 1:1-1:20 tagged
create stpd s8
enable stpd s8 auto-bind v5
configure stpd s8 tag 100
enable stpd s8
```

To automatically add ports to the CIST STPD and to expand the boundary of the STPD, you must complete the following tasks:

- Create a VLAN or use the Default VLAN. (In this example, the Default VLAN is used.)
- Create the MSTP region.
- Create the STPD to be used as the CIST, and configure the mode of operation for the STPD.
- Specify the priority for the CIST.
- Enable the CIST.

The following example enables autobind on the VLAN Default for the CIST STPD named *s1*:

```
configure mstp region 1
create stpd s1
configure stpd s1 mode mstp cist
configure stpd s1 priority 32768
enable stpd s1
```

The following example enables autobind on the VLAN math for the MSTI STPD named *s2*:

```
create vlan math
configure vlan math tag 2
configure vlan math add ports 2-3
configure mstp region 1
create stpd s2
configure stpd s2 mode mstp msti 1
configure stpd s2 priority 32768
enable stpd s2 auto-bind vlan math
configure stpd s2 ports link-type point-to-point 5-6
enable stpd s2
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd ports

```
enable stpd <stpd_name> ports [all | <port_list>]
```

Description

Enables the STP protocol on one or more ports.

Syntax Description

stpd_name	Specifies an STPD on the switch.
all	Specifies all ports for a given STPD.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

If STP is enabled for a port, BPDUs are generated and processed on that port if STP is enabled for the associated STPD.

You must configure one or more STPDs before you can use the `enable stpd ports` command. To create an STPD, use the `create stpd <stpd_name>` command. If you have considerable knowledge and experience with STP, you can configure the STPD using the `configure stpd` commands. However, the default STP parameters are adequate for most networks.

Example

The following command enables slot 2, port 4 on an STPD named *Backbone_st*:

```
enable stpd backbone_st ports 2:4
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd rapid-root-failover

```
enable stpd <stpd_name> rapid-root-failover
```

Description

Enables rapid root failover for faster STP recovery times.

Syntax Description

stpd_name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

If you create an STPD with a unique name, the keyword `stpd` is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command enables rapid root fail over on STPD *Backbone_st*:

```
enable stpd backbone_st rapid-root-failover
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show stpd

```
show stpd {<stpd_name> | detail}
```

Description

Displays STPD settings on the switch.

Syntax Description

stpd_name	Specifies an STPD on the switch.
detail	Specifies that STPD settings should be shown for each STPD.

Default

N/A.

Usage Guidelines

If you specify the command without any options, the following STPD information appears:

- Name—The name of the STPD.
- Tag—The StpdID of the domain, if configured.
- Flags—The following flags communicate information about the current state of the STPD:
 - (C) Topology Change—A network topology change has occurred in the network.
 - (D) Disable—The STPD is disabled.
 - (E) Enable—The STPD is enabled.
 - (R) Rapid Root Failover—The STPD has been configured for rapid root failover
 - (T) Topology Change Detected—The STPD has detected a change in the network topology.
 - (M) MSTP CIST—The STPD has been configured for MSTP, and the STPD is the common and internal spanning tree.
 - (I) MSTP MSTI—The STPD has been configured for MSTP, and the STPD is a multiple instance spanning tree.
- Ports—The number of ports that are part of the STPD.
- Bridge ID—The MAC addresses of the switch.
- Designated Root—The MAC address of the switch that is the designated root bridge.
- Rt Port—The root port.
- Rt Cost—The path cost to the root port.
- Total Number of STPDs—The total number of STPDs configured on the switch.

If you have an MSTP region and associated spanning trees configured on the switch, the command also displays the following global MSTP information:

- MSTP Region—The name of the MSTP region configured on the switch.
- Format Identifier—The number used by BPDUs to communicate within an MSTP region.

- Revision Level—This number is reserved for future use.
- Common and Internal Spanning Tree (CIST)—The name of the CIST that controls the connectivity of interconnecting MSTP regions.
- Total number of MST Instances (MSTI)—The number of MSTIs running in the MSTP region.

If you use the `show stpd` command and specify the name of an STPD, in addition to the data previously described, the command displays more detailed information about the STPD. If you specify the `detail` option, the switch displays the same type of information for all of the STPDs configured on the switch.

The additional output includes the following:

- STPD mode of operation
- Autobind mode
- Active VLANs
- Timer information
- Topology change information

If you have MSTP configured, the command also displays the following information:

- Bridge role
- CIST root
- CIST regional root
- MSTI instances
- Master port (Displayed only on MSTI STPDs)

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test"
```

In this example, to view the settings of the STPD *Test*, enter `show stpd Test`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the STPD settings on a switch that has MSTP configured:

```
show stpd
```

The following is sample output from this command:

```
MSTP Global Configuration:
MSTP Region Name           : 00304841ed97
MSTP format Identifier      : 0
MSTP Revision Level        : 3
Common and Internal Spanning Tree (CIST) : ----
Total Number of MST Instances (MSTI)    : 0

Name Tag  Flags  Ports Bridge ID          Designated Root  Rt Port Rt Cost
s0  0000D----- 0 8000001030f99dc0 0000000000000000 ----- 0
Total number of STPDs: 1
```

Flags: (C) Topology Change, (D) Disable, (E) Enable, (R) Rapid Root Failover
(T) Topology Change Detected, (M) MSTP CIST, (I) MSTP MSTI

The following command displays STPD settings on an STPD named *Backbone_st*:

```
show stpd backbone_st
```

The following is sample output from this command:

```
Stpd: backbone_st Stp: ENABLED      Number of Ports: 51
Rapid Root Failover: Disabled
Operational Mode: 802.1W           Default Binding Mode: 802.1D
802.1Q Tag: (none)
Ports: 1:1,1:2,2:1,2:2,3:1,3:2,4:1,4:2,5:1,5:2,
       5:3,5:4,5:5,5:6,5:7,5:8,5:9,5:10,5:11,5:12,
       5:13,5:14,5:15,5:16,5:17,5:18,5:19,5:20,5:21,5:22,
       5:23,5:24,5:25,5:26,5:27,5:28,5:29,5:30,5:31,5:32,
       5:33,5:34,5:35,5:36,5:37,5:38,5:39,5:40,5:41,5:42,
       5:43
Participating Vlan: Default
Auto-bind Vlan: Default
Bridge Priority: 5000
BridgeID:          13:88:00:01:30:f4:06:80
Designated root:   0a:be:00:01:30:28:b7:00
RootPathCost: 19    Root Port: 28
MaxAge: 20s         HelloTime: 2s         ForwardDelay: 15s
CfgBrMaxAge: 20s    CfgBrHelloTime: 2s        CfgBrForwardDelay: 15s
Topology Change Time: 35s          Hold time: 1s
Topology Change Detected: FALSE     Topology Change: FALSE
Number of Topology Changes: 7
Time Since Last Topology Change: 4967s
```

The following is sample output for an STPD configured as the CIST (the output is similar for an STPD configured as an MSTI):

```
Stpd: s0              Stp: DISABLED      Number of Ports: 0
Rapid Root Failover: Disabled
Operational Mode:      MSTP              Default Binding Mode: 802.1d
MSTP Instance :        CIST              CIST : s0
802.1Q Tag:            (none)
Ports:                 (none)
Participating Vlan Count: 1
Auto-bind Vlan Count:  1

Bridge Priority:        32768
BridgeID:              80:00:00:10:30:f9:9d:c0Bridge
Role :                 CIST Regional Root
CIST Root              80:00:00:10:30:f9:9d:c0CIST
Regional Root:         80:00:00:10:30:f9:9d:c0
Designated root:       00:00:00:00:00:00:00:00
RootPathCost: 0        External RootPathCost: 0      Root Port: ----
MaxAge: 0s             HelloTime: 0s             ForwardDelay: 0s
CfgBrMaxAge: 20s       CfgBrHelloTime: 2s         CfgBrForwardDelay: 15s
MaxHopCount: 20        CfgBrMaxHopCount : 20

Topology Change Time:   35s                 Hold time: 1s
Topology Change Detected: FALSE             Topology Change: FALSE
```



```
Number of Topology Changes: 0
Since Last Topology Change: 0s

Participating Vlans      : (none)
Auto-bind Vlans          : Default
```

History

This command was available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show stpd ports

```
show stpd <stpd_name> ports {[detail | <port_list> {detail}}]
```

Description

Displays the STP state of a port.

Syntax Description

stpd_name	Specifies an STPD name.
port_list	Specifies one or more ports or slots and ports.
detail	Specifies more detailed information about one or more ports of the STPD.

Default

N/A.

Usage Guidelines

This command displays the following:

- STPD port configuration
- STPD port encapsulation mode
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root designated, alternate and so on)
- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type
- Edge port settings (inconsistent behavior, edge safeguard setting)
- MSTP port role (internal or boundary)

To display more detailed information for one or more ports in the specified STPD, including participating VLANs, specify the `detail` option.

If you have MSTP configured and specify the `detail` option, this command displays additional information:

- MSTP internal path cost
- MSTP timers

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test ports"
```

In this example, to view all of the port settings of STPD *Test*, enter `show stpd Test ports`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the state of slot 3, ports 1 through 3 on an STPD named *s0*:

```
show stpd S0 ports 3:1-3:3
```

The following is sample output from this command:

Port	Mode	State	Cost	Flags	Priority	Port ID	Designated Bridge
3:1	802.1D	FORWARDING	100	e-----	16	16641	00:00:00:00:00:00:00:00
3:2	802.1D	FORWARDING	100	e-----	16	16642	00:00:00:00:00:00:00:00
3:3	802.1D	FORWARDING	100	e-----	16	16643	00:00:00:00:00:00:00:00

Total Ports: 3

```

----- Flags: -----
1:                e=Enable, d=Disable
2: (Port role)    R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type)  b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)   b=broadcast, p=point-to-point, e=edge
5:                p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:                i = edgeport inconsistency
8:                S = edgeport safe guard active
                  s = edgeport safe guard configured but inactive
9:                B = Boundary, I = Internal

```

The following command displays the detailed information for the ports in STPD *S0*:

```
show stpd S0 ports detail
```

The following is sample output from this command:

```

Stpd: s0          Port: 1:1          PortId: 8002      Stp: ENABLED      Path Cost: 19
External Path Cost:
Port Mode: 802.1D
Port State: FORWARDING          Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  00:00:00:00:00:00:00:00      Designated Cost: 0
Designated Bridge: 00:00:00:00:00:00:00:00      Designated Port Id: 0
Partner STP version: Dot1d
Edge Port Safe Guard: Disabled
Participating Vlans: Default

Stpd: s0          Port: 1:2          PortId: 8002      Stp: ENABLED      Path Cost: 19
External Path Cost:
Port Mode: 802.1D

```

```

Port State: FORWARDING           Topology Change Ack: FALSE
Port Priority: 16
Designated Root: 00:00:00:00:00:00:00:00      Designated Cost: 0
Designated Bridge: 00:00:00:00:00:00:00:00     Designated Port Id: 0
Partner STP version: Dot1d
Edge Port Safe Guard: Disabled
Participating Vlans: Default

```

The following command displays the detailed information for the ports in STPD *s1* configured for MSTP:

```
show stpd s1 ports detail
```

The following is sample output from this command:

```

Stpd: s1          Port: 1 PortId: 8001      Stp: ENABLED      Path Cost: 4
Port Mode: 802.1D
Port State: FORWARDING           Topology Change Ack: FALSE
Port Priority: 16
Designated Root: 80:00:00:04:96:1f:a8:44      Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44     Designated Port Id: 8001
Partner STP version: MSTP
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15      helloTime: 2      maxHops: 20
Participating Vlans: v1

Stpd: s1          Port: 2 PortId: 8002      Stp: ENABLED      Path Cost: 4
Port Mode: 802.1D
Port State: BLOCKING             Topology Change Ack: FALSE
Port Priority: 16
Designated Root: 80:00:00:04:96:1f:a8:44      Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44     Designated Port Id: 8002
Partner STP version: Dot1d
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15      helloTime: 2      maxHops: 20
Participating Vlans: v1

```

History

This command was available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show vlan stpd

```
show vlan <vlan_name> stpd
```

Description

Displays the STP configuration of the ports assigned to a specific VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If you have a VLAN that spans multiple STPDs, use this command to display the STP configuration of the ports assigned to that specific VLAN.

This command displays the following:

- STPD port configuration
- STPD port mode of operation
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root designated, alternate and so on)
- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type

If your VLAN has the same name as another component, for example an STPD, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the `vlan` keyword, the switch displays an error message similar to the following:

```
%% Ambiguous command: "show Test stpd"
```

In this example, to view the STPD state of VLAN *Test*, enter `show vlan Test stpd`.

If you enter a VLAN name that is not associated with an STPD or does not exist, the switch displays an error message similar to the following:

```
Failed to find vlan 'vlan1' or it has no STP domains configured on it
```

If this happens, check to make sure you typed the correct name of the VLAN and that the VLAN is associated with an STPD.

If your VLAN has a name unique only to that VLAN, the keyword `vlan` is optional.

Example

The following command displays the spanning tree configurations for the vlan *Default*:

```
show vlan default stpd
```

The following is sample output from this command:

```
s0(enabled)  Tag: (none)  Ports: 8 Root/P/C: 80:00:00:01:30:94:79:00/-----/0
```

Port	Mode	State	Cost	Flags	Priority	Port ID	Designated Bridge
1:1	802.1D	LEARNING	19	eDbb-d-	16	8001	80:00:00:01:30:94:79:00
1:2	802.1D	DISABLED	4	e-----	16	8002	00:00:00:00:00:00:00:00
1:3	802.1D	DISABLED	4	e-----	16	8003	00:00:00:00:00:00:00:00
1:4	802.1D	LEARNING	4	eDbb-d-	16	8004	80:00:00:01:30:94:79:00
1:5	802.1D	LEARNING	4	eDbb-d-	16	8005	80:00:00:01:30:94:79:00
1:6	802.1D	DISABLED	4	e-----	16	8006	00:00:00:00:00:00:00:00
1:7	802.1D	DISABLED	4	e-----	16	8007	00:00:00:00:00:00:00:00
1:8	802.1D	DISABLED	4	e-----	16	8008	00:00:00:00:00:00:00:00

```

----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role)  R=Root, D=Designated, A=Alternate, B=Backup, M=Master, Y=Boundary
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)  b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d=802.1d, w=802.1w, m=mstp
7:          i=edgeport inconsistency
8:          B = Boundary, I = Internal

```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure mstp region

```
unconfigure mstp region
```

Description

Unconfigures the MSTP region on the switch and returns all MSTP settings to their default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Before you unconfigure an MSTP region, Extreme Networks recommends that you disable all active STPDs in the region. This includes the CIST and any active MSTIs.

After you issue this command, all of the MSTP settings return to their default values, as described below:

- **Region Name**—This indicates the name of the MSTP region. In the Extreme Networks implementation, the maximum length of the name is 32 characters and can be a combination of alphanumeric characters and underscores (_).
- **Format Selector**—This indicates a number to identify the format of MSTP BPDUs. The default is 0.
- **Revision Level**—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command unconfigures the MSTP region on the switch:

```
unconfigure mstp region
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

unconfigure stpd

```
unconfigure stpd {<stpd_name>}
```

Description

Restores default STP values to a particular STPD or all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.

Default

N/A.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

Use this command to restore default STP values to a particular STPD. If you want to restore default STP values on all STPDs, do not specify a spanning tree name.

Example

The following command restores default values to an STPD named *Backbone_st*:

```
unconfigure stpd backbone_st
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure stpd ports link-type

```
unconfigure stpd <stpd_name> ports link-type <port_list>
```

Description

Returns the specified port to the factory default setting of broadcast link.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

All ports are broadcast link types.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, you must enter the `stpd` keyword to specify the STPD. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Example

The following command configures slot 2, ports 1 through 4 to return to the factory default of broadcast links in STPD `s1`:

```
unconfigure stpd s1 ports link-type 2:1-2:4
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

This chapter describes the commands for:

- Enabling and disabling ESRP
- Performing ESRP configuration
- Enabling and disabling port restart and failure tracking for ESRP
- Displaying ESRP configuration information
- Enabling and disabling ELRP in an ESRP environment

Extreme Standby Router Protocol

The Extreme Standby Router Protocol (ESRP) is a feature of ExtremeXOS that allows multiple switches to provide redundant Layer 3 routing services to users. In addition to providing Layer 3 routing redundancy, ESRP also provides for Layer 2 redundancy. You can use these “layered” redundancy features in combination or independently. The Layer 2 redundancy features of ESRP offer fast failure recovery and provide for dual-homed system design. In some instances, depending on network system design, ESRP can provide better resiliency than using the Spanning Tree Protocol (STP) or Virtual Router Redundancy Protocol (VRRP).

You configure ESRP on a per domain basis on each switch. A maximum of two switches can participate in providing redundant Layer 3 or Layer 2 services to a single virtual LAN (VLAN). If you configure and use ESRP groups, more than two switches can provide redundant Layer 2 or Layer 3 services to a single vlan. The switches exchange keep-alive packets for each VLAN independently. Only one switch (the master) can actively provide Layer 3 routing and/or Layer 2 switching for each VLAN. This switch handles the forwarding, address resolution protocol (ARP) requests, and forwarding for this particular VLAN. Other participating switches for the VLAN are in slave mode waiting for an ESRP state change.

To participate in ESRP, the following must be true:

- A VLAN can belong to only one ESRP domain.
- The IP address for the VLANs participating in an ESRP domain must be identical.
- All switches in the ESRP network must use the same election algorithm, otherwise loss of connectivity, broadcast storms, or other unpredictable behavior may occur.
- If you have an untagged master VLAN, you must specify an ESRP domain ID. The domain ID must be *identical* on all switches participating in ESRP for that particular domain.
- If you have a tagged master VLAN, ESRP uses the 802.1Q tag (VLANid) of the master VLAN for the ESRP domain ID. If you do not use the VLANid as the domain ID, you must specify a different domain ID. As previously described, the domain ID must be identical on all switches participating in ESRP for that particular domain.

Extreme Networks recommends that all switches participating in ESRP run the same version of ExtremeXOS. Not all ESRP features are available in all ExtremeXOS software releases.

Licensing

You must have an Advanced Edge, Core, or Advanced Core license to configure and use all of the ESRP features described in this chapter.

If you have an Edge license, you can configure your Extreme Networks switch to be ESRP-aware. ESRP-aware switches do not actively participate in ESRP; however, when connected to a network with ESRP-enabled switches, the ESRP-aware switches reliably perform failover and failback scenarios in the prescribed recovery times.

The BlackDiamond 10808 switch with an MSM-1 module or an MSM1-XL module ships with a Core or Advanced Core license, respectively.

The BlackDiamond 12804 switch with an MSM-5 module or an MSM-5R module ships with a Core license. The BlackDiamond 12802 switch ships with an Advanced Edge license.

The BlackDiamond 8800 series switch with an MSM-G8X module or an MSM-48 module, Summit X450 series switch, and Summit X450a series switch ship with an Advanced Edge license.

The Summit X450e series switches ship with an Edge license, which only supports ESRP-aware. You can, however, upgrade to an Advanced Edge license that supports ESRP.

SummitStack requires an Advanced Edge license or Core license.

For more information about software licensing, including how to obtain and upgrade your license, see [Chapter 2, “ExtremeXOS Overview,”](#) in the *ExtremeXOS Concepts Guide*.

ESRP Modes of Operation

ExtremeXOS has two modes of ESRP operation: standard and extended. Select standard if your network contains some switches running ExtremeWare, others running ExtremeXOS, and a combination of those switches participating in ESRP. Standard ESRP is backward compatible with and supports the ESRP functionality of ExtremeWare.

Select extended ESRP if your network contains switches running *only* ExtremeXOS. Extended mode ESRP supports and is compatible with switches running ExtremeXOS. By default, ExtremeXOS operates in extended mode.

ESRP-Aware Switches

Extreme Networks switches that are not actively participating in ESRP but are connected on a network that has other Extreme Networks switches running ESRP, are ESRP-*aware*. When ESRP-aware switches are attached to ESRP-enabled switches, the ESRP-aware switches reliably perform failover and failback scenarios in the prescribed recovery times. For an Extreme Networks switch to be ESRP-aware, you must create an ESRP domain on the aware switch (however, you do not enable the ESRP domain), add a master VLAN to that ESRP domain, and configure a domain ID, if necessary.

To participate as an ESRP-aware switch, the following must be true:

- The ESRP domain name must *identical* on all switches (ESRP-enabled and ESRP-aware) participating in ESRP for that particular domain.

- The master VLAN name and IP address must be *identical* on all switches (ESRP-enabled and ESRP-aware) participating in ESRP for that particular domain.
- The domain ID must be *identical* on all switches (ESRP-enabled or ESRP-aware) participating in ESRP for that particular domain.
 - If you have an untagged master VLAN, you must specify an ESRP domain ID.
 - If you have a tagged master VLAN, ESRP uses the 802.1Q tag (VLANid) of the master VLAN for the ESRP domain ID. If you do not use the VLANid as the domain ID, you must specify a different domain ID.

Extreme Loop Recovery Protocol

ExtremeXOS 11.1 introduces support for the Extreme Loop Recovery Protocol (ELRP). ELRP is a feature of ExtremeXOS that allows you to prevent, detect, and recover from Layer 2 loops in the network. You can use ELRP with other protocols such as ESRP.

With ELRP, each switch, except for the sender, treats the ELRP PDU as a Layer 2 multicast packet. The sender uses the source and destination MAC addresses to identify the packet it sends and receives. When the sender receives its original packet back, that triggers loop detection and prevention. Once a loop is detected, the loop recovery agent is notified of the event and takes the necessary actions to recover from the loop. ELRP operates only on the sending switch; therefore, ELRP operates transparently across the network.

Details of using the standalone ELRP client are discussed in the *ExtremeXOS User Guide* chapter, “[Troubleshooting](#),” and the commands used to configure standalone ELRP are described in [Appendix B](#), “[Troubleshooting Commands](#).”

clear esrp counters

```
clear esrp counters
```

Description

Clears the statistics gathered by ESRP for all ESRP domains on the switch.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

Use this command to clear the state transition and the protocol packet counters gathered by ESRP.

The state transition count displays the number of times the ESRP domain entered the following states:

- **Aware**—An Extreme switch that does not participate in ESRP elections but is capable of listening to ESRP Bridge Protocol Data Units (BPDUs).
- **Master**—The master switch is the device with the highest priority based on the election algorithm. The master is responsible for responding to clients for Layer 3 routing and Layer 2 switching for the ESRP domain.
- **Neutral**—The neutral state is the initial state entered by the switch. In a neutral state, the switch waits for ESRP to initialize and run. A neutral switch does not participate in ESRP elections.
- **PreMaster**—The pre-master state is an ESRP switch that is ready to be master but is going through possible loop detection prior to transitioning to master.
- **Slave**—The slave switch participates in ESRP but is not elected or configured the master and does not respond to ARP requests but does exchange ESRP packets with other switches on the same VLAN. The slave switch is available to assume the responsibilities of the master switch if the master becomes unavailable or criteria for ESRP changes.

If the slave is in extended mode, it does not send ESRP hello messages; however, it sends PDUs that may trigger a change in the master switch.

For more information about configuring the ESRP mode of operation on the switch, see the `configure esrp mode [extended | standard]` command. By default, ESRP operates in extended mode.

To display information about the ESRP domain, including the previously described states, use the `show esrp {<name>}` command.

The protocol packet count displays the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

To display information about the ESRP counters, use the `show esrp {<name>} counters` command.

Example

The following command clears the statistics gathered by ESRP:

```
clear esrp counters
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

clear esrp neighbor

```
clear esrp <esrpDomain> neighbor
```

Description

Clears the neighbor information for the specified ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.

Default

N/A.

Usage Guidelines

If you add a new switch to your ESRP domain, use this command to clear the existing neighbor information for the ESRP domain. After the switch is up, running, and configured as an ESRP-aware or ESRP-enabled device, new neighbor information is learned.

Before using this command, schedule a downtime for your network. Use this command for maintenance purposes only.

Example

The following command clears the existing neighbor information on the ESRP domain *esrp1* after adding a new switch to the ESRP domain:

```
clear esrp esrp1 neighbor
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

clear esrp sticky

```
clear esrp <esrpDomain> sticky
```

Description

Clears the stickiness in the ESRP domain and forces the election of the ESRP master switch.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.

Default

N/A.

Usage Guidelines

Use the `clear esrp sticky` command to force the election of the ESRP master switch. Before using this command, schedule a downtime for your network.

For example, without stickiness configured, if an event causes the ESRP master to failover to the backup, the previous backup becomes the new master. If another event causes the new master to return to backup, you have experienced two network interruptions. To prevent this, use the `configure esrp election-policy` and select stickiness as an election algorithm.

If you use sticky as an election metric, and an event causes the ESRP master to failover, ESRP assigns the new master with the highest sticky election metric of 1. Therefore, regardless of changes to the neighbor's election algorithm, the new master retains its position. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

If you use `clear esrp <esrpDomain> sticky` command, it only affects the current master and can trigger ESRP re-election.

Example

The following command clears the stickiness on the ESRP domain *esrp1*:

```
clear esrp esrp1 sticky
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add elrp-poll ports

```
configure esrp <esrpDomain> add elrp-poll ports [<ports> | all]
```

Description

Configures the ports of an ESRP domain where ELRP packet transmission is requested by ESRP.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ports	Specifies list of slots and ports.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

This command allows you to configure the ports in your network that might experience loops, such as ports that connect to master, slave, or ESRP-aware switches, to receive ELRP packets. You do not need to send ELRP packets to host ports.

Example

The following command enables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain *esrp1*:

```
configure esrp esrp1 add elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp add master

```
configure esrp <esrpDomain> add master <vlan_name>
```

Description

Adds a master VLAN to an ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.
vlan_name	Specifies the name of the master VLAN.

Default

N/A.

Usage Guidelines

You must configure one *master* VLAN for each ESRP domain. A master VLAN can belong to one ESRP domain only. An ESRP domain contains one master and zero or more member VLANs.

The master VLAN:

- Exchanges ESRP PDUs, hello messages, and data between a pair of ESRP-enabled switches.
- Contains the total number of active physical ports that are counted when determining the master ESRP domain. The switch with the highest number of active ports takes priority.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds the master VLAN *purple* to the ESRP domain *esrp1*:

```
configure esrp esrp1 add master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add member

```
configure esrp <esrpDomain> add member <vlan_name>
```

Description

Adds a member VLAN to an ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.
vlan_name	Specifies the name of the member VLAN.

Default

N/A.

Usage Guidelines

You can configure zero or more *member* VLANs for each ESRP domain. An ESRP domain contains one master and zero or more member VLANs.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds the member VLAN *green* to the ESRP domain *esrp1*:

```
configure esrp esrp1 add member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-environment

```
configure esrp <esrpDomain> add track-environment failover <priority>
```

Description

Configures an ESRP domain to track environmental failures.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
priority	Specifies a number between 0 and 254. The default priority is 255. See the following "Usage Guidelines" section for more information.

Default

No environmental tracking.

Usage Guidelines

Environmental tracking tracks power supply and chassis temperature status.

If a failure is detected, the ESRP domain priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the domain, it will cause the affected domain to go into slave mode.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch will remain in slave mode even when the VLAN fails over from the current master.

To make effective use of this feature, the normal priority of the ESRP domain must be higher than the failover priority of this command.

Example

The following command enables environmental failure tracking, and specifies that the ESRP priority for ESRP domain *esrp1* be set to 10 upon an environmental failure.

```
configure esrp esrp1 add track-environment failover 10
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-iproute

```
configure esrp <esrpDomain> add track-iproute <ipaddress>/<masklength>
```

Description

Configures an ESRP domain to track a route entry in the system's routing table.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ipaddress	Specifies the IP address of the route entry to be tracked.
masklength	Specifies the subnet of the route entry to be tracked.

Default

Disabled.

Usage Guidelines

The track-ip metric consists of the total number of tracked IP routes that are “up” or functional.

An ESRP domain can track eight IP routes.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
configure esrp esrp1 add track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-ping

```
configure esrp <esrpDomain> add track-ping <ipaddress> frequency <seconds>
miss <misses>
```

Description

Configures an ESRP domain to track an external gateway using ping.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ipaddress	Specifies the IP address of the external gateway.
seconds	Specifies the interval in seconds between ping requests. The range is 1 to 600 seconds.
misses	Specifies the number of consecutive ping failures that will initiate failover to an ESRP slave. The range is 1 to 256 pings.

Default

No ping tracking.

Usage Guidelines

The tracked-ping metric consists of the total number of stations that are successfully tracked using ping. ESRP uses an aggregate of tracked pings and traced routes to track an external gateway.

An ESRP domain can track eight stations.

Example

The following command enables ping tracking for the external gateway at 10.207.29.17, pinging every 10 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure esrp esrp1 add track-ping 10.207.29.17 frequency 10 miss 5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-vlan

```
configure esrp <esrpDomain> add track-vlan <vlan_name>
```

Description

Configures an ESRP domain to track port connectivity to a specified VLAN.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
vlan_name	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

The track-vlan metric is derived from the total number of active physical ports on the VLAN being tracked by the ESRP domain.

If more than one VLAN shares a physical link, each VLAN counts the physical link.

The ESRP switch should have a higher priority number than its neighbors to ensure master election.

An ESRP domain can track one VLAN, and the tracked VLAN should not be a member of any other ESRP domain in the system.

Example

The following command enables ESRP domain *esrp1* to track port connectivity to VLAN *engineering*:

```
configure esrp esrp1 add track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp aware add selective-forward-ports

```
configure esrp <domain> aware add selective-forward-ports <portlist> {group  
<group number>}
```

Description

Enables selective forwarding by creating an aware port list and adds additional ports to the list.

Syntax Description

domain	Specifies an ESRP domain name.
portlist	Specifies the ports to be added to the aware port list.
group number	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware VLAN. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic.

Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This ESRP-aware port list is then used for forwarding ESRP PDUs.

Use this command to create or add to an existing port list for the ESRP groups associated with an ESRP-aware VLAN.

Example

The following command configures esrp domain (d1) to forward ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware add selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on all platforms.

configure esrp aware delete selective-forward-ports

```
configure esrp <domain> aware delete selective-forward-ports <all|portlist>
{group <group number>}
```

Description

Disables all or part of selective forwarding by deleting ports from the ESRP-aware port list.

Syntax Description

domain	Specifies an ESRP domain name.
all	Specifies that all of the ports are to be disabled.
portlist	Specifies the ports to be disabled from the ESRP-aware port list.
group number	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

By configuring selective forwarding, you create an ESRP-aware port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. That port list is used for forwarding ESRP PDUs from the selected ports only of an ESRP-aware switch.

Use this command to delete one or more or all of the ports from an ESRP-aware port list. Deleting all of the ports puts the domain back to the default state.

Example

The following command configures esrp domain (d1) to exclude ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware delete selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on all platforms.

configure esrp delete elrp-poll ports

```
configure esrp <esrpDomain> delete elrp-poll ports [<ports> | all]
```

Description

Disables ELRP packet transmission on ports of an ESRP domain.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ports	Specifies list of slots and ports in the ESRP domain.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

If you have host ports on an ESRP domain, you do not need to send ELRP packets to those ports.

If you change your network configuration, and a port no longer connects to a master, slave, or ESRP-aware switch, you can disable ELRP transmission on that port.

Example

The following command disables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain *esrp1*:

```
configure vlan esrp1 delete elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp delete master

```
configure esrp <esrpDomain> delete master <vlan_name>
```

Description

Deletes the specifies master VLAN from the specified ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.
vlan_name	Specifies the name of the master VLAN.

Default

N/A.

Usage Guidelines

You must disable the ESRP domain before removing the master VLAN. To disable the ESRP domain, use the `disable esrp {<esrpDomain>}` command.

If you attempt to remove the master VLAN before disabling the ESRP domain, the switch displays an error message similar to the following:

```
ERROR: Failed to delete master vlan for domain "esrp1" ; ESRP is enabled!
```

If this happens, disable the ESRP domain and re-issue the `configure esrp delete master` command.

Example

The following command deletes the master VLAN *purple* from the ESRP domain *esrp1*:

```
configure esrp esrp1 delete master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete member

```
configure esrp <esrpDomain> delete member <vlan_name>
```

Description

Deletes a member VLAN from the specified ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.
vlan_name	Specifies the name of the member VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the member VLAN *green* from the ESRP domain *esrp1*:

```
configure esrp esrp1 delete member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-environment

```
configure esrp <esrpDomain> delete track-environment
```

Description

Disables environmental failure tracking for an ESRP domain.

Syntax Description

esrpDomain	Specifies an ESRP domain name.

Default

No environmental tracking.

Usage Guidelines

None.

Example

The following command disables environmental failure tracking for ESRP domain *esrp1*:

```
configure esrp esrp1 delete track-environment
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-iproute

```
configure esrp <esrpDomain> delete track-iproute <ipaddress>/<masklength>
```

Description

Disables route entry tracking for an ESRP domain.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ipaddress	Specifies the IP address of the route entry to be tracked.
masklength	Specifies the subnet of the route entry to be tracked.

Default

Disabled.

Usage Guidelines

If you disable route tracking for a failed route, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a route that is up and functional, there is no impact on the ESRP state.

Example

The following command disables tracking of routes to the specified subnet for ESRP domain *esrp1*:

```
configure esrp esrp1 delete track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-ping

```
configure esrp <esrpDomain> delete track-ping <ipaddress>
```

Description

Disables the tracking of an external gateway using ping.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ipaddress	Specifies the IP address of the external gateway.

Default

No ping tracking.

Usage Guidelines

If you disable ping tracking for a failed ping, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a successful ping, there is no impact on the ESRP state.

Example

The following command disables ping tracking for the external gateway at 10.207.29.17:

```
configure esrp esrp1 delete track-ping 10.207.29.17
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-vlan

```
configure esrp <esrpDomain> delete track-vlan <vlan_name>
```

Description

Disables the tracking of port connectivity to a specified VLAN.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
vlan_name	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

If you delete a VLAN that is down, the ESRP domain recovers from the forced standby state.

Example

The following command disables the tracking of port connectivity to VLAN *engineering*:

```
configure esrp esrp1 delete track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp domain-id

```
configure esrp <esrpDomain> domain-id <number>
```

Description

Assigns an ESRP domain ID to an ESRP domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.
number	Specifies the number to use for the ESRP domain ID. The user-configured ID range is 4096 through 65,535.

Default

If the master VLAN is tagged, ESRP uses that VLANid for the ESRP domain ID. If the master VLAN is untagged, you must specify the ESRP domain ID.

Usage Guidelines

Before you enable a specific ESRP domain, it must have a domain ID. A domain ID is either a user-configured number or the VLANid of the tagged master VLAN. If you do not have a domain ID, you cannot enable ESRP on that domain.

Each switch participating in ESRP for a particular domain must have the same domain ID configured.

The `number` parameter range for user-configured domain IDs is 4096 through 65,535.

If the master VLAN is tagged, you can use that VLANid for the ESRP domain ID. The range for VLAN tags is 2 through 4095. Tag 1 is assigned to the default VLAN.

Example

The following command assigns the domain ID 5000 to ESRP domain *esrp1*:

```
configure esrp esrp1 domain-id 5000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp election-policy

```
configure esrp <esrpDomain> election-policy [ports > track > priority
| ports > track > priority > mac | priority > mac
| priority > ports > track > mac | priority > track > ports > mac
| sticky > ports > track > priority
| sticky > ports > track > priority > mac
| sticky > ports > weight > track > priority > mac
| sticky > priority > mac
| sticky > priority > ports > track > mac
| sticky > priority > track > ports > mac
| sticky > track > ports > priority
| sticky > track > ports > priority > mac
| track > ports > priority | track > ports > priority > mac]
```

Description

Configures the election algorithm on the switch.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority.
ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority, MAC address. NOTE: This is the default election algorithm for standard mode.
priority > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, MAC address.
priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, active ports, tracking information, MAC address.
priority > track > ports > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, tracking information, active ports, MAC address.
sticky > ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority.
sticky > ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority, MAC address.
sticky > ports > weight > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, port weight, tracking information, ESRP priority, MAC address. NOTE: Beginning with ExtremeXOS 11.1 and later, this is the default election algorithm for extended mode.
sticky > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, MAC address.
sticky > priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, active ports, tracking information, MAC address.

<code>sticky > priority > track > ports > mac</code>	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, tracking information, active ports, MAC address.
<code>sticky > track > ports > priority</code>	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority.
<code>sticky > track > ports > priority > mac</code>	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority, MAC address.
<code>track > ports > priority</code>	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority.
<code>track > ports > priority > mac</code>	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority, MAC address.

Default

In extended mode, the default election algorithm is `sticky > ports > weight > track > priority > mac`.

In standard mode, the default election algorithm is `ports > track > priority > mac`.

Usage Guidelines

The election algorithm determines the order of precedence of the election factors used to determine the ESRP Master. The election factors are:

- Stickiness (`sticky`): the switch with the higher sticky value has higher priority. When an ESRP domain claims master, its sticky value is set to 1 (available in extended mode only).
- Active Ports (`ports`): the number of active ports (the switch with the highest number takes priority)
- Tracking Information (`track`): whether the switch is using ESRP tracking. A switch using tracking has priority.
- ESRP Priority (`priority`): a user-defined priority number between 0 and 254. A higher number has higher priority. The default priority setting is 0. A priority setting of 255 makes an ESRP switch a standby switch that remains in slave mode until you change the priority setting. Extreme Networks recommends this setting for system maintenance. A switch with a priority setting of 255 will never become the master.
- MAC address (`mac`): the switch MAC address. A higher-number address has priority.
- Active port weight (`weight`)—The switch that has the highest port weight takes precedence. The bandwidth of the port automatically determines the port weight (available only in extended mode). ESRP does not count ports with a weight of 0 (known as don't count ports) regardless of ESRP running in extended or standard mode.

The election algorithm must be the same on all switches for a particular ESRP domain. The election algorithms that use `sticky` are and `weight` are available in extended mode only.

In ExtremeXOS 11.0, the extended mode default election algorithm is: `sticky > ports > track > priority > mac > weight`. This election algorithm is not supported in ExtremeXOS 11.1.

Factors to Consider. The `ports-track-priority` or `track-ports-priority` options can be used to ensure that there is no failback if the original Master recovers (the Master will have the same ports, tracks and priority, but a higher MAC).

Any of the options with sticky can also be used to ensure that there is no failback if the original master recovers. With sticky, if an event causes the ESRP master to failover, ESRP assigns the new master with the sticky count of 1. After sticky is set on the master, regardless of changes to its neighbor's election algorithm, the new master retains its position. For example, adding active ports to the slave does not cause the new master to failback to the original master, even if the slave has more active ports than the master. Sticky algorithms provide for fewer network interruptions than non-sticky algorithms. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

Switch Behavior . If a switch is master, it actively provides Layer 3 routing services to other VLANs, and Layer 2 switching between all the ports of that VLAN. Additionally, the switch exchanges ESRP packets with other switches that are in slave mode.

If a switch is in slave mode, it exchanges ESRP packets with other switches on that same VLAN. When a switch is in slave mode, it does not perform Layer 3 routing or Layer 2 switching services for the VLAN.

Updating the Election Algorithm . ESRP uses the default election policy for extended mode. If you have an ESRP domain operating in standard mode, the domain ignores the sticky and weight algorithms. To change the election algorithm, you must first disable the ESRP domain and then configure the new election algorithm. If you attempt to change the election algorithm without disabling the domain first, an error message appears.

To disable the ESRP domain, use the following command:

```
disable esrp {<esrpDomain>}
```

To modify the election algorithm, use the following command:

```
configure esrp <esrpDomain> election-policy [ports > track > priority | ports > track
> priority > mac | priority > mac | priority > ports > track > mac | priority > track
> ports > mac | sticky > ports > track > priority | sticky > ports > track > priority
> mac | sticky > ports > weight > track > priority > mac | sticky > priority > mac |
sticky > priority > ports > track > mac | sticky > priority > track > ports > mac |
sticky > track > ports > priority | sticky > track > ports > priority > mac | track >
ports > priority | track > ports > priority > mac]
```

If you attempt to use an election algorithm not supported by the switch, an error message similar to the following appears:

```
ERROR: Specified election-policy is not supported!
```

Supported Policies:

1. sticky > ports > weight > track > priority > mac
2. ports > track > priority
3. sticky > ports > track > priority
4. ports > track > priority > mac
5. sticky > ports > track > priority > mac
6. priority > mac
7. sticky > priority > mac
8. priority > ports > track > mac
9. sticky > priority > ports > track > mac
10. priority > track > ports > mac
11. sticky > priority > track > ports > mac

```
12. track > ports > priority
13. sticky > track > ports > priority
14. track > ports > priority > mac
15. sticky > track > ports > priority > mac
```

Example

The following command configures the election algorithm to use tracking information as the first criteria for determining the ESRP master switch for ESRP domain *esrp1*:

```
configure esrp esrp1 election-policy track > ports > priority > mac
```

History

This command was first available in ExtremeXOS 11.0.

The default election algorithm for extended mode was updated to `sticky > ports > weight > track > priority > mac`, and the `weight` election factor was used in ExtremeXOS 11.1. The `sticky > ports > track > priority > mac > weight` election algorithm is not supported in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-master-poll disable

```
configure esrp <esrpDomain> elrp-master-poll disable
```

Description

Disables the use of ELRP by ESRP in the master state.

Syntax Description

esrpDomain	Specifies an ESRP domain name.

Default

Disabled.

Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the master state. When you disable ELRP, the ESRP master switch no longer transmits ELRP PDUs to detect network loops.

Example

The following command disables the use of ELRP in the master state on ESRP domain *elrp1*:

```
configure esrp elrp1 esrp elrp-master poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-master-poll enable

```
configure esrp <esrpDomain> elrp-master-poll enable {interval <interval>}
```

Description

Enables the use of ELRP by ESRP in the master state, and configures how often the master checks for loops in the network.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
interval	Specifies how often, in seconds, successive ELRP packets are sent. The default is 1 second. The range is 1 to 64 seconds.

Default

- Use of ELRP in the master state—disabled
- Interval—1 second

Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the master state. When an ESRP-enabled switch is in the master state, and you enable `elrp-master-poll`, the switch periodically sends ELRP PDUs at the configured interval level. If a loop is detected in the network, the transmitted PDUs are received by the switch. The ESRP master switch then transitions to the slave state to break the network loop.

Extreme Networks recommends that you enable both premaster and master polling when using ELRP with ESRP. To enable premaster polling, use the `configure esrp <esrpDomain> elrp-premaster-poll enable {count <count> | interval <interval>}`.

If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

```
ERROR: Premaster-poll should be enabled before enabling master-poll!
```

If this happens, first configure premaster polling followed by master polling (if required).

Specify the `interval` parameter to configure how often successive ELRP PDUs are sent while in the master state. If you do not specify an interval value, the default value is used.

Example

The following command enables the use of ELRP in the master state on ESRP domain `elrp1`:

```
configure esrp elrp1 esrp elrp-master poll enable
```

The following command configures the ESRP master to check for loops in the network every 3 seconds:

```
configure esrp elrp1 esrp elrp-master-poll enable interval 3
```


History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-premaster-poll disable

```
configure esrp <esrpDomain> elrp-premaster-poll disable
```

Description

Disables the use of ELRP by ESRP in the pre-master state.

Syntax Description

esrpDomain	Specifies an ESRP domain name.

Default

Disabled.

Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the pre-master state. When you disable ELRP in the pre-master state, the ESRP pre-master switch no longer transmits ELRP PDUs to detect network loops prior to changing to the master state.

Example

The following command disables the use of ELRP in the pre-master state on the ESRP domain *elrp1*:

```
configure esrp elrp1 esrp elrp-premaster poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-premaster-poll enable

```
configure esrp <esrpDomain> elrp-premaster-poll enable {count <count> |
interval <interval>}
```

Description

Enables the use of ELRP by ESRP in the pre-master state, and configures how many times the switch sends ELRP PDUs and how often the switch sends ELRP PDUS in the pre-master state.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
count	Specifies the number of times the switch sends ELRP PDUs. The default is 3. The range is 1 to 32.
interval	Specifies how often, in seconds, the ELRP PDUs are sent. The default is 1 second. The range is 1 to 32 seconds.

Default

- Use of ELRP in the pre-master state—disabled
- Count—3 times
- Interval—1 second

Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the pre-master state to prevent network loops from occurring. When an ESRP-enabled switch is in the pre-master state (waiting to become the master), and you enable elrp-premaster-poll, the switch periodically sends ELRP PDUs at the configure level for a specified number of times. If there is a loop in the network, the transmitted PDUs are received by the switch. If this happens, the ESRP pre-master switch does not transition to the master state; rather, the switch transitions to the slave state.

Extreme Networks recommends that you enable both premaster and master polling when using ELRP with ESRP. To enable master polling, use the `configure esrp <esrpDomain> elrp-master-poll enable {interval <interval>}`.

If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

```
ERROR: Premaster-poll should be enabled before enabling master-poll!
```

If this happens, first configure premaster polling followed by master polling (if required).

If you do not specify the optional `count` or `interval` parameters, the default values are used.

If the sender does not receive packets, there is no loop in the network.

Example

The following command enables the use of ELRP—with the default settings—in the pre-master state on ESRP domain *elrp1*:

```
configure esrp elrp1 esrp elrp-premaster poll enable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp group

```
configure esrp <esrpDomain> group <number>
```

Description

Configures the group number to be used for the ESRP domain.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
group_number	Specifies the ESRP group number to which this ESRP domain should be added. The range is 0 through 31.

Default

The default group number is 0.

Usage Guidelines

Each group runs an instance of ESRP within the same VLAN or broadcast domain. A maximum of seven ESRP groups can be defined within the same networked broadcast domain. In addition, a maximum of seven distinct ESRP groups can be supported on a single ESRP switch. You can configure a maximum of 32 ESRP groups in a network.

The range for the `group_number` parameter is 0 through 31.

The most typical application for multiple ESRP groups is when two or more sets of ESRP switches are providing fast-failover protection within a common subnet for two or more groups of users. An additional use for ESRP groups is ESRP Host Attach; ESRP VLANs that share the same ESRP HA ports must be members of different ESRP groups.

You must first disable an ESRP domain before you modify an existing or add a new group number. If you try to modify the group number without disabling the ESRP domain, an error message similar to the following is displayed:

```
ERROR: can't change ESRP group for active domain "esrp1"!
```

To disable an ESRP domain, use the `disable esrp {<esrpDomain>}` command.

Example

The following command configures ESRP domain `esrp1` to be a member of ESRP group 2:

```
configure esrp esrp-1 group 2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp mode

```
configure esrp mode [extended | standard]
```

Description

Configures the mode of operation for ESRP on the switch.

Syntax Description

extended	Specifies ESRP extended mode. This mode is compatible with devices running ExtremeXOS and is not backward compatible with devices running ExtremeWare.
standard	Specifies ESRP standard mode. This mode is backward compatible with devices running ExtremeWare.

Default

The default mode is extended.

Usage Guidelines

Use standard ESRP if your network contains a combination of switches running ExtremeXOS and ExtremeWare participating in ESRP. With standard ESRP, the switches running ExtremeXOS are backward compatible with the switches running ExtremeWare.

Use extended ESRP if your network contains switches running *only* ExtremeXOS; this is the default.

If your network has switches currently running ExtremeWare, and you add a BlackDiamond 10808 switch, a BlackDiamond 8800 series switch, a BlackDiamond 12800 series switch, SummitStack, or a Summit family switch running ExtremeXOS, select standard ESRP. By selecting standard, the switch running ExtremeXOS is backward compatible with the ExtremeWare implementation of ESRP.

If you use the default mode—extended—and your ESRP domain contains a switch running ExtremeXOS that detects a neighbor switch running ExtremeWare, the mode automatically changes to standard for that domain. This action causes the switch to enter the neutral state and re-elect the ESRP master. Since you are using the default mode of operation, and the switch running ExtremeXOS detected a neighbor switch running ExtremeWare, the ExtremeXOS switch toggles to standard mode although the configured mode of operation remains as extended.

Example

The following command configures ESRP to run in standard mode:

```
configure esrp mode standard
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp name

```
configure esrp <esrpDomain> name <new-name>
```

Description

Renames an existing ESRP domain.

Syntax Description

esrpDomain	Specifies the current name of an ESRP domain.
new-name	Specifies a new name for the ESRP domain.

Default

N/A.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, esrp. Names must start with an alphabetical character, for example, a, Z.

You can rename an ESRP domain regardless of its current state.

Example

The following command renames ESRP domain *esrp1* to *esrp3*:

```
configure esrp esrp1 name esrp3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports mode

```
configure esrp ports <ports> mode [host | normal]
```

Description

Configures the ESRP port mode for ESRP host attach.

Syntax Description

ports	Specifies one or more ports or slots and ports that should be configured.
host	Specifies that the ports should be configured as host ports.
normal	Specifies that the ports should be configured as normal ports.

Default

The default port mode is normal.

Usage Guidelines

Ports configured as normal ports do not accept or transmit Layer 2 or Layer 3 traffic when the local ESRP device is a slave.

Ports configured as host ports allow the network to continue operation independent of ESRP status. The command sets the port to forward, allowing those ports directly attached to the slave's hosts to communicate with other hosts that are connected to the master. If you use load sharing with the ESRP HA feature, configure the load-sharing group first and then enable Host Attach on the group.

A Layer 2 connection for VLANs between ESRP switches is required.

An ESRP Host Attach port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command configures ports 1 through 5 on slot 3 as host ports:

```
configure esrp port 3:1-3:5 mode host
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports no-restart

```
configure esrp ports <ports> no-restart
```

Description

Disables port restart for a port.

Syntax Description

ports	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables port restart for ports 7-9 in slot 3 in the ESRP master domain:

```
configure esrp port 3:7-3:9 no-restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports restart

```
configure esrp ports <ports> restart
```

Description

Configures ESRP to restart ports if there is a state change and the downstream switch is from another vendor.

Syntax Description

ports	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

If an ESRP domain becomes a slave, ESRP disconnects member ports that have port restart enabled. The disconnection of these ports causes downstream devices to remove the ports from their FDB tables.

After 3 seconds the ports re-establish connection with the ESRP-enabled device. This feature allows you to use ESRP in networks that include equipment from other vendors.

If switch becomes a slave, ESRP disconnects the physical links of member ports that have port restart enabled.

An ESRP restart port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command enables port restart for ports 7-9 in slot 3 on the ESRP master domain:

```
configure esrp port 3:7-3:9 restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports weight

```
configure esrp ports <ports> weight [auto | <port-weight>]
```

Description

Assigns the port weight for the specified ESRP port(s).

Syntax Description

ports	Specifies one or more ports or slots and ports.
auto	Specifies the switch to calculate the weight of a port based on the port's bandwidth and link speed.
port-weight	Specifies an ESRP port weight of 0. With a port weight of 0, the ports are not counted.

Default

The switch automatically calculates the weight of a port based on the bandwidth of the port.

Usage Guidelines

Use this command to override the automatically calculated port weight.

The `port-weight` parameter specifies a weight of 0. With this configuration, ESRP does not count host ports and normal ports as active. With a weight of 0, ESRP experiences fewer state changes due to frequent client activities like rebooting and unplugging laptops. A don't-count port cannot be a mirroring, software-controlled redundant port, or a Netlogin port.

For load shared ports, configure one master port in the load-share group with the port weight. A single command specifies the weight for the entire load shared group. You can specify any port from the load share group in the command. A load-shared port has an aggregate weight of all of its member ports. If you add or delete a member port (or trunk), the weight of the master load-shared port is updated. For more information about load sharing, see [“Configuring Slots and Ports on a Switch”](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command configures port 1 on slot 3 with a weight of 0:

```
configure esrp port 3:1 weight 0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp priority

```
configure esrp <esrpDomain> priority <number>
```

Description

Configures the ESRP priority.

Syntax Description

esrpDomain	Specifies an ESRP domain number.
number	Specifies a number between 0 and 255.

Default

The default ESRP priority is 0.

Usage Guidelines

The ESRP priority is one of the factors used by the ESRP election algorithm in determining which switch is the Master switch.

The range of the priority value is 0 to 254, with 0 being the lowest priority, 254 being the highest. If the ESRP priority is the determining criteria for the election algorithm, the highest priority value determines which switch will act as master for a particular ESRP domain.

Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch will remain in slave mode even when the ESRP domain fails over from the current master. This feature is typically used to ensure a switch cannot become the ESRP master while it is offline for servicing.

Example

The following command configures the ESRP priority to the highest priority on ESRP domain *esrp1*:

```
configure esrp esrp1 priority 254
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer hello

```
configure esrp <esrpDomain> timer hello <seconds>
```

Description

Configures the ESRP hello timer value.

Syntax Description

esrpDomain	Specifies an ESRP domain name
seconds	Specifies the number of seconds between keep-alive packets. The range is 1 to 255 seconds.

Default

The default hello timer is 2 seconds.

Usage Guidelines

The timer specifies the interval, in seconds, for exchanging keep-alive packets between the ESRP switches for this ESRP domain. A lower value specifies a more frequent exchange of keep-alive messages, resulting in the faster detection of a failover condition. The timer setting must be configured identically for the ESRP domain across all participating switches. To see the hello settings, use the `show esrp {<name>}` command.

The `seconds` range is 1 to 255.

If your configuration contains more than 2,000 ESRP VLANs and 256,000 FDB entries, Extreme Networks recommends a timer setting greater than 3 seconds.

To view the hello timer settings, use the `show esrp {<name>}` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when FDB entries are flushed during a master-slave transition. To avoid this we recommend the general neighbor and hello timeout guidelines listed in [Table 25](#) described in the `configure esrp timer neighbor` command.

Example

The following command configures the ESRP hello timer to 4 seconds for the ESRP domain *esrp1*:

```
configure esrp esrp1 timer hello 4
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer neighbor

```
configure esrp <esrpDomain> timer neighbor <seconds>
```

Description

Configures the ESRP neighbor timeout value.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
seconds	Specifies the number of seconds after which an ESRP neighbor times out. The range is 6 to 1024 seconds.

Default

The default neighbor timeout is 8 seconds (four times the hello timer).

Usage Guidelines

The neighbor timeout specifies the amount of time that ESRP waits before considering the neighbor down. The neighbor value must be at least 3 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ; neutral
timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello to 1024 seconds.

To view the neighbor timer settings, use the `show esrp {<name>}` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when FDB entries are flushed during a master-slave transition. To avoid this Extreme Networks recommends the general neighbor and hello timeout guidelines listed in [Table 25](#).

Table 25: General neighbor and hello timeout

Number of Domains	Number of VLANs	Suggested Neighbor and Hello Timeout
64 or less	1000	Use the default timer values
64	1000 to 3000	hello >=3, neighbor >=9
128	3000	hello >=4, neighbor >=12

Example

The following command configures the ESRP neighbor timeout to 14 seconds for the ESRP domain *esrp1*:

```
configure esrp esrp1 timer neighbor 14
```


History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer neutral

```
configure esrp <esrpDomain> timer neutral <seconds>
```

Description

Configures the ESRP neutral timeout value.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
seconds	Specifies the number of seconds after which an ESRP domain. The range is 4 to 1024 seconds.

Default

The default neutral timeout is 4 seconds (two times the hello timer).

Usage Guidelines

After you create, configure, and enable the ESRP domain, it enters the neutral state. The neutral timeout specifies the amount of time the ESRP domain stays in this temporary state before entering the slave state. The neutral value must be at least 2 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ; neutral
timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 2*hello to 1024.

To view the neutral timer settings, use the `show esrp {<name>}` command.

Example

The following command configures the ESRP neutral timeout to 8 seconds for the ESRP domain *esrp1*:

```
configure esrp esrp1 timer neutral 8
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer premaster

```
configure esrp <esrpDomain> timer premaster <seconds>
```

Description

Configures the ESRP pre-master timeout value.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
seconds	Specifies the maximum length of time, in seconds, that the transitioning master VLAN remains in the pre-master state. The range is 6 to 1024.

Default

The default timeout is 6 seconds (three times the hello timer).

Usage Guidelines

The premaster timer specifies how long the ESRP domain stays in the pre-master state. The pre-master timer expires if the neighbor agrees to be the slave. The premaster value must be at least three times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ; neutral
timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello-1024.

To view the pre-master timer settings, use the `show esrp {<name>}` command.



CAUTION

Configure the pre-master state timeout only with guidance from Extreme Networks personnel. Misconfiguration can severely degrade the performance of ESRP and your switch.

Example

The following command configures the pre-master timeout to 10 seconds for the ESRP domain `esrp1`:

```
configure esrp esrp-1 timer premaster 10
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer restart

```
configure esrp <esrpDomain> timer restart <seconds>
```

Description

Configures the ESRP restart timer value.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
seconds	Specifies the maximum length of time, in seconds, that the neighbor ESRP switch remains in its current state during an MSM hitless failover. The range is 15 to 1024.

Default

The default restart timer value is 30 seconds.

Usage Guidelines

The restart timer specifies the amount of time that the neighbor ESRP switch remains in its current state during a hitless failover. This timer prevent the slave ESRP switch from trying to become master during a hitless failover.

The `seconds` range is 15-1024.

To view the restart settings, use the `show esrp {<name>}` command.

Example

The following command configures the restart timer value to 40 seconds for the ESRP domain *esrp1*:

```
configure esrp esrp-1 timer restart 40
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create esrp

```
create esrp <esrpDomain>
```

Description

Creates an ESRP domain with the specified name on the switch.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain to be created. May be up to 32 characters in length.

Default

The ESRP domain is disabled and in the “Aware” state.

When you create an ESRP domain, it has the following default parameters:

- Operational version—Extended
- Priority—0
- VLAN interface—none
- VLAN tag—0
- Hello timer—2 seconds
- Neighbor timer—8 seconds
- Premaster timer—6 seconds
- Neutral timer—4 seconds
- Neighbor restart timer—30 seconds
- VLAN tracking—none
- Ping tracking—none
- IP route tracking—none

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, esrp. Names must start with an alphabetical character, for example, a, Z.

Each ESRP domain name must be unique and cannot duplicate any other named ESRP domains on the switch. If you are uncertain about the ESRP names on the switch, use the `show esrp` command to view the ESRP domain names.

You can create a maximum of 128 ESRP domains.

Configuring ESRP-Aware Switches. For an Extreme Networks switch to be ESRP-aware, you must create an ESRP domain on the aware switch, add a master VLAN to that ESRP domain, add a member VLAN to that ESRP domain if configured, and configure a domain ID if necessary.

On switches with an Edge license, only ESRP-aware functionality is supported; you cannot enable the ESRP domain. On switches with an Advanced Edge, a Core, or an Advanced Core license, you can configure switches as ESRP-aware or with full ESRP functionality. If you need an ESRP-aware switch, do not enable the ESRP domain.

For more information about configuring ESRP-aware switches, see the [Chapter 21, “Extreme Standby Router Protocol,”](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command creates ESRP domain *esrp1* on the switch:

```
create esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

delete esrp

```
delete esrp <esrpDomain>
```

Description

Deletes the ESRP domain with the specified name.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain to be deleted.

Default

N/A.

Usage Guidelines

You must first disable an ESRP domain before you delete it. To disable an ESRP domain, use the `disable esrp` command.

You do not have to remove the master or member VLANs from an ESRP domain before you delete it. When you delete an ESRP domain, All VLANs are automatically removed from the domain.

Example

The following command deletes ESRP domain *esrp1* from the switch:

```
delete esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable esrp

```
disable esrp {<esrpDomain>}
```

Description

Disables ESRP for a named domain or for the entire switch.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.

Default

Disabled for the entire switch.

Usage Guidelines

If you do not specify a domain name, ESRP is disabled for the entire switch.

If you disable an ESRP domain, the switch notifies its neighbor that the ESRP domain is going down, and the neighbor clears its neighbor table. If the master switch receives this information, it enters the neutral state to prevent a network loop. If the slave switch receives this information, it enters the neutral state.

Example

The following command disables ESRP for the entire switch:

```
disable esrp
```

The following command disables ESRP for the domain *esrp1*:

```
disable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable esrp

```
enable esrp <esrpDomain>
```

Description

Enables ESRP for a named domain.

Syntax Description

esrpDomain	Specifies the name of an ESRP domain.

Default

Disabled.

Usage Guidelines

Before you enable an ESRP domain, it must have a domain ID. The ESRP domain ID is determined from one of the following user-configured parameters:

- ESRP domain number created with the `configure esrp domain-id` command
- 802.1Q tag (VLANid) of the tagged master VLAN

If you do not have a domain ID, you cannot enable ESRP on that domain. A message similar to the following appears:

```
ERROR: Cannot enable ESRP Domain "esrp1" ; No domain id configured!
```

If you add an untagged master VLAN to the ESRP domain, make sure to create an ESRP domain ID with the `configure esrp domain-id` command before you attempt to enable the domain.

Example

The following command enables ESRP for the domain *esrp1*:

```
enable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms with an Advanced Edge, a Core, or an Advanced Core license.

show esrp

```
show esrp {<name>}
```

Description

Displays ESRP configuration information for one or all ESRP domains on the switch.

Syntax Description

name	Specifies an ESRP domain name.

Default

Shows summary ESRP information.

Usage Guidelines

This command shows information about the state of an ESRP domain and its neighbors. This includes information about tracked devices.

In addition to ESRP information, ELRP status information is also displayed. This includes information about the master and pre-master states, number of transitions to the pre-master state, and the ports where ELRP is disabled.

The output varies depending upon the configuration and the state of the switch.

Example

The following command displays summary ESRP status information for the ESRP domains on the switch:

```
show esrp
```

The following is sample output from this command:

```
ESRP:                               Enabled
Configured Version:                 Extended
# ESRP domain configuration :
-----
Domain Grp Ver VLAN   VID  DId  IP/IPX      State  Master MAC Address Nbr
-----
ed2     0  E   v2     2    2   2.2.2.3     Master 00:01:30:f9:9e:90  0
ed2     5  E   v2     2    2   2.2.2.3     Aware  00:01:00:0D:9e:8a  0
ed3     0  E   v3     3    3   0.0.0.0     Aware  00:01:00:0C:F0:D1  0
ed4     0  E   v4     4    4   0.0.0.0     Slave  00:00:00:00:00:00  0
-----
```

ESRP Port configuration:

```
-----
Port      Weight      Host      Restart
-----
6:1       0             H
6:2       10
6:3       0             R
-----
```

The following command displays detailed ESRP status information for the specified ESRP domain on the switch (the election policy displayed is the default policy in extended mode):

```
show esrp ed2
```

The following is sample output from this command:

```
show esrp ed2
Domain:                ed2
Group:                 0
Operational Version:   extended
Vlan Interface:        v2
Vlan Tag:              2
Domain Id:             2
Rtif. Admin Status:    DOWN
Rtif. Virtual Mac :    00:e0:2b:00:00:80
IP Address:            2.2.2.3
Election Policy:
standby > sticky > ports > weight > track > priority > mac
-----
                                This System                Neighbor system
-----
State:                        Slave
Active Ports:                 1
Tracked Active Vlan Ports:    0
Tracked IP Routes & Pings:    0
Priority:                     255[255]
Sticky Flag:                  0
MAC:                          00:01:30:f9:9e:90
Active Ports Weight:          10
Sequence Number:              8
Hand Shake Flag:              0
Restart Flag:                 0
Timer Configuration:          Hello                2s(0)      Neighbor      6s(0)
PreMaster 4s(0) Neutral 4s(0)
NbrRestart 30s(0)
State Transition Counters: To Master 0 To Neutral 1
To PreMaster 0 To Aware 1
To Slave 1
Last State Change: Mon Apr  5 10:43:34 2004

ELRP in Premaster:           Enabled (Interval: 1, Count: 3)
ELRP in Master:              Enabled (Interval: 1)

Tracked Vlans: t_vlan
Tracked Pings: 40.0.1.2 / 3 second intervals / 5 misses allowed
Tracked Ip Routes: 30.4.2.16/255.255.255.0
Tracked Environment: System Power : OK
```

Temperature : WARNING

Domain Member VLANs:

VLAN Name	VID	Virtual IP/IPX	State
m_vlan1	1001	0.0.0.0	DOWN
m_vlan2	1002	0.0.0.0	DOWN
m_vlan3	1003	0.0.0.0	DOWN

History

This command was first available in ExtremeXOS 11.0.

ELRP status information was added in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show esrp aware

```
show esrp <domain> aware
```

Description

Displays all selective forwarding information for a given ESRP-aware domain.

Syntax Description

domain	Specifies the name of an ESRP domain.

Default

Disabled

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware VLAN. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic. Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This aware port list is then used for forwarding ESRP PDUs.

Use this command to display all ESRP-aware information for a given domain. You can also specify either

- Selective Forward Ports only with the command: `show esrp <domain> aware selective-forward-ports` or
- Statistics only with the command: `show esrp <domain> aware statistics`.

Example

The following command displays the ESRP aware information for the domain d1.

```
show esrp d1 aware
```

The command produces output similar to the following two tables:

```
Domain: d1
Vlan:   vesrp1
-----
Group   Port Count   Selective Forward Ports
-----
0       5             5:1, 5:2, 7:31, 7:32: 8:1
3       2             5:1, 8:1
```

Group	Master MAC	Rx	Fwd	FDB Flush	Fwd Ports
0	00:12:00:33:44:55	10	10	1	selective
1	00:22:00:12:21:1F	77	77	3	default
3	00:02:00:13:11:11	99	99	3	selective

The first of the two tables displays Selective Forward Ports information:

- **Group**—The number of an ESRP group within the given domain.
- **Port Count**—The number of ports in the group that are selected to forward PDUs on the master VLAN.
- **Selective Forward Ports**—The list of ports in the group that are selected to forward PDUs on the master VLAN.

The second of the two tables displays statistical information about the activity of the ports:

- **Group**—The number of an ESRP group within the given domain
- **Master MAC**—The MAC address for the master of the group.
- **Rx**—The number of PDUs received matching the domain/group pair.
- **Fwd**—The number of PDUs received and forwarded matching the domain/group pair.
- **FDB Flush**—The number of FDB Flush requests received from the ESRP Master for this domain/group pair.
- **Fwd Ports**—Selective or Default.

Selective describes the group as having a configured aware port list for selective forwarding of PDUs on the Master VLAN. The list of ports is displayed in the first table above.

Default describes the group as one where all the ports on the master VLAN forward the ESRP PDUs that are received for the domain/group pair. Because there is no selective forwarding configuration for this group, there is no entry in the first table.

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on all platforms.

show esrp counters

```
show esrp {<name>} counters
```

Description

Displays ESRP counter information for ESRP or for a specified ESRP domain.

Syntax Description

name	Specifies an ESRP domain name.

Default

Displays summary ESRP counter information.

Usage Guidelines

The `show esrp counters` command displays information about the number of:

- Failed received protocol packets
- Failed sent protocol packets
- Dropped protocol packets belonging to unknown ESRP domains
- Dropped protocol packets due to invalid Extreme Encapsulation Protocol (EEP) data
- Dropped packets due to old sequence numbers
- Dropped packets due to an invalid 802.1Q tag
- Dropped packets because the packet length was truncated (packet length is less than expected)
- Dropped packets due to failed checksum verification

The `show esrp {<name>} counters` command displays information about the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

Example

The following command displays ESRP counter information:

```
show esrp counters
```

The following is sample output from this command:

```
Current-time:                Sun Nov 16 00:25:08 2003
esrpStatsRxHelloFailed       = 0
esrpStatsTxHelloFailed       = 0
esrpStatsUnknownDomain       = 0
esrpStatsUnsupportedEEPVersion = 0
esrpStatsInvalidEEPLength    = 0
esrpStatsNotInTimeWindow     = 0
esrpStatsInvalid8021Qtag     = 0
```

```
esrpStatsInvalidSNAPType      = 0
esrpStatsUndersizePkt         = 0
esrpStatsInvalidChecksum      = 0
esrpStatsWrongDigest          = 0
```

The following command displays counter information for ESRP domain ed5:

```
show esrp ed5 counters
```

The following is sample output from this command:

```
Domain:  ed5                Current-time:  Sun Nov 16 00:25:27 2003
Rx-Esrp-Pkts      = 628      Tx-Esrp-Pkts      = 630
Rx-Aware-Esrp-Pkts = 112     Tx-Aware-Esrp-Pkts = 34
Rx-Err-Pkts       = 0        Tx-Err-Pkts       = 0
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

This chapter describes commands for:

- Enabling and disabling Virtual Router Redundancy Protocol (VRRP)
- Performing basic VRRP configuration

Like Extreme Standby Router Protocol (ESRP), VRRP is a protocol that allows multiple switches to provide redundant routing services to users. A virtual router is a group of one or more physical devices that acts as the default gateway for hosts on the network. The virtual router is identified by a virtual router identifier (VRID) and an IP address. All of the VRRP routers that participate in the virtual router are assigned the same VRID.

The Extreme Networks implementation of VRRP implementation is compliant with RFC 2338, Virtual Router Redundancy Protocol.

The following points pertain to VRRP:

- VRRP packets are encapsulated IP packets.
- The VRRP multicast address is 224.0.0.18.
- The virtual router MAC address is 00 00 5E 00 01 <vrid>
- An interconnect link between VRRP routers should not be used, except when VRRP routers have hosts directly attached.
- A maximum of 128 VRID instances are supported on the router.
- Up to 7 unique VRIDs can be configured on an interface. VRIDs can be re-used, but not on the same interface.
- VRRP and Spanning Tree can be simultaneously enabled on the same switch.
- VRRP and ESRP cannot be simultaneously enabled on the same VLAN.

VRRP uses an election algorithm to dynamically assign responsibility for the master router to one of the VRRP routers on the network. A VRRP router is elected master if one of the following is true, the router is:

- The IP address owner (router that has the IP address of the virtual router configured as its real interface address).
- Configured with the highest priority (the range is 1 - 255).

In the event of a tie in priority, the highest primary IP address has precedence.

If the master router becomes unavailable, the election process provides dynamic failover and the backup router that has the highest priority assumes the role of master.

A new master is elected when one of the following things happen:

- VRRP is disabled on the master router.
- Communication is lost between master and backup router(s). The master router sends periodic advertisements to the backup routers to indicate that it is alive.

VRRP also supports the following tracking options:

- VRRP VLAN tracking
- VRRP route table tracking
- VRRP ping tracking

If a tracking option is enabled, and the object being tracked becomes unreachable, the master device will fail over.

Licensing

You must have an Advanced Edge, Core, or Advanced Core license to configure and use all of the VRRP features described in this chapter.

VRRP functionality is not available with an Edge license.

The BlackDiamond 10808 switch with an MSM-1 module or an MSM1-XL module ships with a Core or Advanced Core license, respectively.

The BlackDiamond 12804 switch with an MSM-5 module or an MSM-5R module ships with a Core license. The BlackDiamond 12802 switch ships with an Advanced Edge license.

The BlackDiamond 8800 series switch with an MSM-G8X module or an MSM-48 module, the Summit X450 and X450a series switches ship with an Advanced Edge license.

The Summit X450e series switches ship with an Edge license, which does not support VRRP. You can, however, upgrade to an Advanced Edge license that supports VRRP.

A SummitStack must be operating at the Advanced Edge level or greater to support all VRRP features.

For more information about software licensing, including how to obtain and upgrade your license, see [Chapter 2, “ExtremeXOS Overview,”](#) in the *ExtremeXOS Concepts Guide*.

clear counters vrrp

```
clear counters vrrp {{vlan <vlan_name>} {vrid <vridval>}}
```

Description

Clears, resets all VRRP statistics and counters.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

Use this command to reset the VRRP statistics on the switch. Statistics are not reset when you disable and re-enable VRRP.

If you do not enter a parameter, statistics for all VRRP VLANs are cleared.

If you specify only VLAN name, statistics for all VRRP VRIDs on that VLAN are cleared.

If you specify VLAN name and VRRP VRID, only statistics for that particular VRID are cleared.

Example

The following command clears the VRRP statistics on VRRP VLAN *v1*:

```
clear counters vrrp vlan v1
```

The following command clears the VRRP statistics for VRID *1* on VRRP VLAN *v1*:

```
clear counters vrrp vlan v1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid add ipaddress

```
configure vrrp vlan <vlan_name> vrid <vridval> add <ipaddress>
```

Description

Associates a virtual IP address with a specific VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
ipaddress	Specifies the IP address of the virtual router in which this device participates.

Usage Guidelines

The restrictions on this command are as follows:

- If the priority of the VRRP virtual router is 255, the IP address to be added must be owned by the VLAN on which the VRRP virtual router exists. If the priority is not 255, the IP address must not be owned by that VLAN.
- When a VRRP virtual router is enabled, it must have at least one virtual IP address.

Example

Create a VLAN named *vlan-1* with an ipaddress of 10.1.2.2 and a VRRP VRID of 1:

```
create vlan vlan-1
configure vlan vlan-1 ipaddress 10.1.2.2
create vrrp vlan-1 vrid 1
```

The following example associates IP address 10.1.2.3 to VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 add 10.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid add track-iproute

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-iproute  
<ipaddress>/<masklength>
```

Description

Creates a tracking entry for the specified route. When this route becomes unreachable, this entry is considered to be failing.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the virtual router ID of the target VRRP virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the prefix of the route to track.
masklength	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

The route specified in this command may not exist in the IP routing table. When you create the entry for a route, an immediate VRRP failover may occur.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
configure vrrp vlan vlan-1 vrid 1 add track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid add track-ping

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-ping <ipaddress>
frequency <seconds> miss <misses>
```

Description

Creates a tracking entry for the specified IP address. The entry is tracked via pings to the IP address, sent at the specified frequency.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the IP address to be tracked.
seconds	Specifies the number of seconds between pings to the target IP address. The range is 1 to 600 seconds.
misses	Specifies the number of misses allowed before this entry is considered to be failing. The range is 1 to 255 pings.

Default

N/A.

Usage Guidelines

Adding an entry with the same IP address as an existing entry will cause the new values to overwrite the existing entry's frequency and miss number.

Example

The following command enables ping tracking for the external gateway at 3.1.0.1, pinging every 3 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure vrrp vlan vlan-1 vrid 1 add track-ping 3.1.0.1 frequency 3 miss 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid add track-vlan

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-vlan
<target_vlan_name>
```

Description

Configures a VRRP VLAN to track port connectivity to a specified VLAN. When this VLAN is in the “down” state, this entry is considered to be failing.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
target_vlan_name	Specifies the name of the VLAN to track.

Default

N/A.

Usage Guidelines

Up to eight VLANs can be tracked.

Deleting a tracked VLAN does not constitute a failover event for the VR tracking it, and the tracking entry will be deleted.

Example

The following command enables VRRP VLAN *vlan-1* to track port connectivity to VLAN *vlan-2*:

```
configure vrrp vlan vlan-1 vrid 1 add track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid advertisement-interval

```
configure vrrp vlan <vlan_name> vrid <vridval> advertisement-interval
<interval> [{seconds} | milliseconds]
```

Description

Configures the time between VRRP advertisements (pings) in seconds or milliseconds.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
interval	Specifies the time interval between advertisements in seconds unless otherwise specified as milliseconds. The default is 1 second. The range is 1 through 255 seconds or 100 through 999 milliseconds. NOTE: If you must configure the range in milliseconds, specify the milliseconds keyword. If you enter a number from 100 through 255 and do not specify the milliseconds keyword, the interval defaults to seconds.

Default

The advertisement interval is 1 second.

Usage Guidelines

The advertisement interval specifies the interval between advertisements sent by the master router to inform the backup routers that its alive. You must use whole integers when configuring the advertisement interval.

An extremely busy CPU can create a short dual master situation. To avoid this, increase the advertisement interval.

The seconds range is 1 through 255. The switch uses the seconds interval by default. You do not need to specify the seconds keyword to configure the advertisement interval in seconds.

The milliseconds range is 100 through 999. Since the switch uses the seconds interval by default, you must specify the milliseconds keyword to configure the advertisement interval in milliseconds.

To view your VRRP configuration, including the configured advertisement interval, use one of the following commands:

- `show vrrp {detail}`
- `show vrrp vlan <vlan_name> {stats}`

Out of Range Error Messages. If you enter a number that is out of the seconds or milliseconds range, the switch displays an error message.

In the following example, the configured advertisement interval is set to 999; however, the `milliseconds` keyword is missing. Since the switch uses the seconds interval by default, and 999 is out of the seconds range, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 250 advertisement-interval 999
Error: Advertisement interval must be between 1 and 255 seconds. 999 out of range
```

If you configure an out of range milliseconds interval, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 2 advertisement-interval 1000 milliseconds
Error: Advertisement interval must be between 100 and 999 milliseconds. 1000 out of range.
```

Example

The following command configures the advertisement interval of 15 seconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 15
```

The following command configures the advertisement interval of 200 milliseconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement -interval 200 milliseconds
```

History

This command was first available in ExtremeXOS 10.1.

The `milliseconds` and `seconds` keywords were added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid authentication

```
configure vrrp vlan <vlan_name> vrid <vridval> authentication [none |  
simplepassword <password>]
```

Description

Configures the authentication type for a specific VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
password	Specifies the user-defined password for authentication.

Default

Authentication is set to *none*.

Usage Guidelines

This command can add a modest amount of security to VRRP advertisements. All VRRP routers using the same VRID must use the same password when using this feature.

A simple password must be between 1 and 8 characters long.

Example

The following command configures authentication for VRRP VLAN *vrrp-1* with the password *newvrrp*:

```
configure vrrp vlan vrrp-1 vrid 1 authentication simple-password newvrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid delete ipaddress

```
configure vrrp vlan <vlan_name> vrid <vridval> delete <ipaddress>
```

Description

Deletes a virtual IP address from a specific VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
ipaddress	Specifies the IP address of the virtual router in which this device participates.

Usage Guidelines

The restrictions on this command are as follows:

- If the priority of the VR is 255, the IP address to be added must be owned by the VLAN on which the VR exists. If the priority is not 255, the IP address must not be owned by that VLAN.
- When a VR is enabled, it must have at least one virtual IP address. When the VR is not enabled, there are no restrictions on deleting the IP address.

Example

The following command removes IP address 10.1.2.3 from VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 delete 10.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid delete track-iproute

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-iproute  
<ipaddress>/<masklength>
```

Description

Deletes a tracking entry for the specified route.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the prefix of the route to track.
masklength	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables tracking of routes to the specified subnet for VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 delete track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid delete track-ping

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-ping  
<ipaddress>
```

Description

Deletes a tracking entry for the specified IP address.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the IP address to be tracked.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled will cause the VRRP VRs state to be re-evaluated for failover.

A VRRP node with a priority of 255 may not recover from a ping-tracking failure if there is a Layer 2 switch between it and another VRRP node. In cases where a Layer 2 switch is used to connect VRRP nodes, Extreme Networks recommends that those nodes have priorities of less than 255.

Example

The following command disables ping tracking for the external gateway at 3.1.0.1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-ping 3.1.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid delete track-vlan

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-vlan  
<target_vlan_name>
```

Description

Deletes the tracking of port connectivity to a specified VLAN.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
target_vlan_name	Specifies the name of the tracked VLAN.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled will cause the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables the tracking of port connectivity to VLAN *vlan-2*:

```
configure vrrp vlan vlan-1 vrid 1 delete track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid dont-preempt

```
configure vrrp vlan <vlan_name> vrid <vridval> dont-preempt
```

Description

Specifies that a higher priority backup router does not preempt a lower priority master.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

The default setting is preempt.

Usage Guidelines

The preempt mode controls whether a higher priority backup router preempts a lower priority master. `dont-preempt` prohibits preemption. The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

Example

The following command disallows preemption:

```
configure vrrp vlan vlan-1 vrid 1 dont-preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid preempt

```
configure vrrp vlan <vlan_name> vrid <vridval> preempt
```

Description

Specifies that a higher priority backup router preempts a lower priority master.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

The default setting is preempt.

Usage Guidelines

The preempt mode controls whether a higher priority backup router preempts a lower priority master. `preempt` allows preemption. The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

Example

The following command allows preemption:

```
configure vrrp vlan vlan-1 vrid 1 preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vrrp vlan vrid priority

```
configure vrrp vlan <vlan_name> vrid <vridval> priority <priorityval>
```

Description

Configures the priority value of a VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
priorityval	Specifies the priority value of the router. The default is 100. The priority range is 1-255.

Default

The default priority is 100.

Usage Guidelines

This command changes the priority of a VRRP VR. If the VR is assigned an IP address that is physically owned by the switch, the VR's priority is 255 and cannot be changed. If the IP address is not owned by switch, the priority cannot be changed to 255.

To change the priority in either of the described cases, disable VRRP and remove the virtual IP address(es) first.

To disable VRRP, use the `disable vrrp {vlan <vlan_name> vrid <vridval>}` command. To remove the virtual IP address(es), use the `configure vrrp vlan <vlan_name> vrid <vridval> delete <ipaddress>` command.

Example

The following command configures a priority of 150 for VLAN *vrrp-1*:

```
configure vrrp vlan vrrp-1 vrid 1 priority 150
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create vrrp vlan vrid

```
create vrrp vlan <vlan_name> vrid <vridval>
```

Description

Creates a VRRP virtual router on the switch.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

VRRP Virtual Router IDs can be used across multiple VLANs. You can create multiple virtual routers on different VLANs. Virtual Router IDs need not be unique to a specific VLAN.

Before configuring any virtual router parameters, you must first create the VRRP instance on the switch. If you define VRRP parameters before creating the VRRP, you may see an error similar to the following:

```
Error: VRRP VR for vlan vrrp1, vrid 1 does not exist.  
Please create the VRRP VR before assigning parameters.  
Configuration failed on backup MSM, command execution aborted!
```

If this happens, create the VRRP instance and then configure its parameters.

Example

The following command creates a VRRP router on VLAN *vrrp-1*, with a VRRP virtual router ID of 1:

```
create vrrp vlan vrrp-1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete vrrp vlan vrid

```
delete vrrp vlan <vlan_name> vrid <vridval>
```

Description

Deletes a specified VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VRRP virtual router on the VLAN *vrrp-1* identified by VRID 2:

```
delete vrrp vlan vrrp-1 vrid 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable vrrp vrid

```
disable vrrp {vlan <vlan_name> vrid <vridval>}
```

Description

Disables a specific VRRP virtual router or all VRRP virtual routers.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

This disables a specific virtual router on the switch. If no VRRP VLAN is specified, all virtual routers on the switch are disabled.

Example

The following command disables all VRRP virtual routers on the switch:

```
disable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable vrrp vrid

```
enable vrrp {vlan <vlan_name> vrid <vridval>}
```

Description

Enables a specific VRRP virtual router or all VRRP virtual routers on the switch.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

This enables a specific virtual router on the device. If you do not specify a VRRP virtual router, all VRRP virtual routers on this device are enabled.

Example

The following command enables all VRRP virtual routers on the switch:

```
enable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vrrp

```
show vrrp {detail}
```

Description

Displays VRRP configuration information for all VRRP VLANs.

Syntax Description

detail	Specifies more detailed VRRP information.

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different VRRP information might be displayed.

If you specify the command without the `detail` keyword, the following VRRP information appears:

- VLAN Name—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows:
 - En—VRRP is enabled on this VLAN.
 - Ds—VRRP is disabled on this VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- Pri—The priority value of the VRRP VLAN.
- Virtual IP Addr—If configured, the virtual IP address associated with the VRRP VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Master Mac Address—The MAC address of the master VRRP router.
- TP—Indicates the number of tracked pings.
- TR—Indicates the number of tracked routes.
- TV—Indicates the number of tracked VLANs.
- P—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- T—Indicates the configured advertisement interval.

If you specify the command with the `detail` keyword, the switch displays similar information in a bulleted format, including:

- VLAN—The name of the VRRP VLAN.

- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- VRRP—The enabled/disabled state of VRRP on the VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Priority—The priority value of the VRRP VLAN.
- Advertisement Interval—Indicates the configured advertisement interval.
- Preempt—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- Authentication—If configured, identifies the VRRP simple password.
- Virtual IP Addresses—If configured, the virtual IP address associated with the VRRP VLAN.
- Tracked Pings—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.
 - Number of misses allowed before this entry is considered to be failing.
- Tracked IP Routes—If configured, displays the IP address and subnet mask of the tracked route(s).
- Tracked VLANs—If configured, displays the name of the tracked VLAN(s).

Example

The following command displays a summary of status information for VRRP:

```
show vrrp
```

The following is sample output from this command:

```
VLAN Name VRID Pri Virtual IP Addr State Master Mac Address TP/TR/TV/P/T
v1(En)      0001 255 1.1.1.1          MSTR 00:00:5e:00:01:01 0 0 0 Y 1

En-Enabled, Ds-Disabled, Pri-Priority, T-Advert Timer, P-Preempt
TP-Tracked Pings, TR-Tracked Routes, TV-Tracked VLANs
```

There are three variations of the detail command that are displayed depending on your switch model.

1 The following command displays detailed status information for VRRP on all platforms:

```
show vrrp detail
```

The following is sample output from this command:

```
VLAN: v1          VRID: 1          VRRP: Enabled State: MASTER
Virtual Router: VR-Default
Priority: 255(master) Advertisement Interval: 1 sec
Preempt: Yes Authentication: simple-password key: foo
Virtual IP Addresses: 1.1.1.1
Tracked Pings: 10.10.10.3 / 10 second intervals / 5 misses allowed
Tracked IP Routes: -
Tracked VLANs: green, purple
* indicates a tracking condition has failed
```

- 2 For a BlackDiamond 10808 or the BlackDiamond 12800 series switch, the same command displays the status information as shown in example #1 but adds the name of the Virtual Router.

```
show vrrp detail
```

The following is sample output from this command on the BlackDiamond 10808 and 12800 switches:

```
VLAN:  vlan_3   VRID:  3           VRRP:  Enabled  State:  BACKUP
Virtual Router:  vir_3
Priority:  100(backup)  Advertisement Interval:  1 sec
Preempt:  Yes   Authentication:  None
Virtual IP Addresses:
  30.1.2.1
Tracked Pings:  -
Tracked IP Routes:  -
Tracked VLANs:  -
```

- 3 When the virtual router is specified, information is displayed only for the domains of that virtual router. Only the BlackDiamond 10808 and the BlackDiamond 12800 series support this feature.

```
show vrrp virtual-router <vr-name> detail
```

The following is sample output from this command on the BlackDiamond 10808 and 12800 switches:

```
VLAN:  vlan_3   VRID:  3           VRRP:  Enabled  State:  MASTER
Virtual Router:  vir_3
Priority:  200(backup)  Advertisement Interval:  1 sec
Preempt:  Yes   Authentication:  None
Virtual IP Addresses:
  30.1.2.1
Tracked Pings:  -
Tracked IP Routes:  -
Tracked VLANs:  -
```

History

This command was first available in ExtremeXOS 10.1.

The Virtual Router name was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show vrrp virtual-router

```
show vrrp virtual-router
```

Description

Displays the VRRP configuration information for all VRRP instances that belong to all created Virtual Routers.

Default

N/A

Usage Guidelines

This displays the following VRRP information:

- VLAN Name—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows
 - En—VRRP is enabled on this VLAN.
 - DS—VRRP is disabled on this VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- Pri—The priority value of the VRRP VLAN.
- Virtual IP Addr—The virtual IP address associated with the VRRP VLAN.
- State—The current state of the VRRP router. The state appears as follows:
 - INIT—The VRRP router is in the initial state.
 - BKUP—The VRRP router is a backup router.
 - MSTR—The VRRP router is the master router.
- Master Mac Address—The MAC address of the master VRRP router.
- Virtual-Router—The virtual router to which a particular VLAN belongs.

Example

This command produces output similar to the following:

VLAN Name	VRID	Pri	Virtual IP Addr	State	Master Mac Address	Virtual-Router
vlan_3 (En)	0003	200	30.1.2.1	MSTR	00:00:5e:00:01:03	vir_3
vlan_4 (En)	0003	200	40.1.2.1	MSTR	00:00:5e:00:01:03	vir_4
vlan_5 (En)	0005	100	50.1.2.1	BKUP	00:00:5e:00:01:05	vir_3
vlan_6 (En)	0005	100	60.1.2.1	BKUP	00:00:5e:00:01:05	vir_4

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and the BlackDiamond 12800 series.

show vrrp virtual-router

```
show vrrp virtual-router <vr-name>
```

Description

Displays the VRRP that matches the specified virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router

Default

N/A

Usage Guidelines

The command displays the VRRP information only for a specific virtual router on the device. If you do not specify a virtual router (refer to `show vrrp virtual-router`), all VRRP virtual routers on this device are displayed.

This displays the following VRRP information:

- VLAN Name—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows:
 - En—VRRP is enabled on this VLAN.
 - Ds—VRRP is disabled on this VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- Pri—The priority value of the VRRP VLAN.
- Virtual IP Addr—If configured, the virtual IP address associated with the VRRP VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Master Mac Address—The MAC address of the master VRRP router.
- TP—Indicates the number of tracked pings.
- TR—Indicates the number of tracked routes.
- TV—Indicates the number of tracked VLANs.
- P—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- T—Indicates the configured advertisement interval.

Example

The following command displays VRRP information for virtual router #3

```
show vrrp vir_3
```

This command produces output similar to the following:

VLAN Name	VRID	Pri	Virtual IP Addr	State	Master Mac Address	TP/TR	TV/P/T
vlan_3(En)	0003	200	30.1.2.1	MSTR	00:00:5e:00:01:03	0 0	0 Y 1
vlan_5(En)	0005	100	50.1.2.1	BKUP	00:00:5e:00:01:05	0 0	0 Y 1

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 10808 and the BlackDiamond 12800 series.

show vrrp vlan

```
show vrrp vlan <vlan_name> {stats}
```

Description

Displays VRRP information for a particular VLAN.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
stats	Specifies statistics for a particular VLAN.

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different VRRP information might be displayed.

If you specify the command without the `stats` keyword, the following VRRP information appears:

- VLAN—The name of the VRRP VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- VRRP—The enabled/disabled state of VRRP on the VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Priority—The priority value of the VRRP VLAN.
- Advertisement Interval—Indicates the configured advertisement interval.
- Preempt—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- Authentication—If configured, identifies the VRRP simple password.
- Virtual IP Addresses—If configured, the virtual IP address associated with the VRRP VLAN.
- Tracked Pings—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.
 - Number of misses allowed before this entry is considered to be failing.
- Tracked IP Routes—If configured, displays the IP address and subnet mask of the tracked route(s).
- Tracked VLANs—If configured, displays the name of the tracked VLAN(s).

If you specify the `stats` keyword, you see counter and statistics information for the specified VRRP VLAN.

Example

The following command displays configuration information for the specified VRRP VLAN:

```
show vrrp vlan blue
```

The following is sample output from this command:

```
VLAN:  blue      VRID:  2      VRRP:  Disabled State:  INIT
Priority:  1(backup)  Advertisement Interval:  1 sec
Preempt:  Yes  Authentication:  None
Virtual IP Addresses:
Tracked Pings:  -
Tracked IP Routes:  -
Tracked VLANs:  -
* indicates a tracking condition has failed
```

The following command displays statistics for VLAN *vrrp-1*:

```
show vrrp vlan vrrp-1 stats
```

The following is sample output from this command:

```
VLAN vrrp-1, VR ID 25
Chksum Err:0, Ver Err:0, VRID Err:0, Auth Mismatch:0, Pkt-len Err:0
Become Master:0, Adv recv:0, Adv Err:0, Auth Fail:0, TTL Err:0
Pri-0-recv:0, Pri-0-snt:0, Addr-List Err:0, Invalid Auth Err:0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

MPLS Commands—BlackDiamond 10808 and 12000 series Switches Only

MPLS encompasses a growing set of protocols defined by the IETF. True to its name, MPLS is based on a label-switching forwarding algorithm. ATM and Frame Relay are examples of other protocols that use label-switching forwarding algorithms.

Conceptually, label switching is straightforward. A label is a relatively short, fixed-length identifier that is used to forward packets received from a given link. The label value is locally significant to a particular link and is assigned by the receiving entity.

Because labels are relatively short (for example, 20 bits in a MPLS shim header), the label of a received packet can be used as an index into a linear array containing the forwarding database. Forwarding database entries indicate the outgoing port and any label(s) to be applied to forwarded frames. Thus, forwarding may consist of a simple lookup and replacement of the incoming label with the appropriate outgoing label (otherwise known as *label swapping*).

This chapter documents the MPLS command set. Some commands are new for MPLS; other commands have been enhanced to support MPLS.

clear counters mpls

```
clear counters mpls {[lsp all | [{vlan} <vlan_name> | vlan all]]}
```

Description

Clears all packet and byte counters for all MPLS LSPs and all MPLS protocol counters for all MPLS interfaces.

Syntax Description

lsp all	Clears all MPLS protocol counters for all MPLS LSPs.
vlan <vlan_name>	Clears all MPLS protocol counters for the MPLS interface on the specified VLAN.
vlan all	Clears all MPLS protocol counters for all MPLS interfaces.

Default

N/A.

Usage Guidelines

This command clears all packet and byte counters for all MPLS LSPs and all MPLS protocol counters for all MPLS interfaces. If the `lsp all` keywords are specified, all packet and byte counters for all MPLS LSPs are cleared. If the `vlan all` keywords are specified, all MPLS protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all MPLS protocol counters for the MPLS interface on that VLAN are cleared.

Example

This example clears all MPLS counters associated with VLAN 1:

```
clear counters mpls vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

clear counters mpls ldp

```
clear counters mpls ldp {{{vlan} <vlan_name>} | lsp all}
```

Description

Clears LDP control protocol counters and packet and byte counters associated with LDP LSPs.

Syntax Description

vlan <vlan_name>	Clears LDP control protocol counters on the specified VLAN.
vlan all	Clears LDP control protocol counters on all MPLS interfaces.
lsp all	Clears all LDP LSP packet and byte counters.

Default

N/A.

Usage Guidelines

By default, all LDP control protocol counters are cleared for all LDP interfaces and all byte counters. Specifying the `vlan` keyword clears only the protocol counters associated with a specified LDP interface. Specifying the `lsp` keyword clears only the packet and byte counters associated with LDP LSPs.

Example

This example clears all LDP control protocol counters and all packet and byte counters for all LDP LSPs:

```
clear counters mpls ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

clear counters mpls rsvp-te

```
clear counters mpls rsvp-te {[lsp all | [{vlan} <vlan_name> | vlan all]]}
```

Description

Clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all MPLS interfaces.

Syntax Description

lsp all	Clears all packet and byte counters for all RSVP-TE LSPs.
vlan <vlan_name>	Clears all RSVP-TE protocol counters for the MPLS interface on the specified VLAN.
vlan all	Clears all RSVP-TE protocol counters on all MPLS interfaces.

Default

By default, all RSVP-TE control protocol counters are cleared for all RSVP-TE interfaces.

Usage Guidelines

This command clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all MPLS interfaces. If the `lsp all` keywords are specified, all packet and byte counters for all RSVP-TE LSPs are cleared. If the `vlan all` keywords are specified, all RSVP-TE protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all RSVP-TE protocol counters for the MPLS interface on that VLAN are cleared.

Example

This example clears the RSVP-TE protocol counters on VLAN 1 only:

```
clear counters mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

clear counters vpls

```
clear counters vpls [<vpls_name> | all]
```

Description

Clears all VPLS counters for the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS VPNs.

Default

N/A.

Usage Guidelines

This command clears all VPLS counters for the specified `vpls_name`. If the optional `all` keyword is specified, all packet and byte counters for all VPLS VPNs are cleared.

Example

This example clears all VPLS counters for the specified VPLS:

```
clear counters vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

clear fdb vpls

```
clear fdb vpls {<vpls_name> {<peer_ip_address>}}
```

Description

Clears the FDB information learned for VPLS.

Syntax Description

<vpls_name>	Clears all FDB entries for the specified VPLS and its associated VLAN.
<peer_ip_address>	Clears all FDB entries for the pseudo wire (PW) associated with the specified VPLS and LDP peer.

Default

N/A.

Usage Guidelines

If the command is used without keywords, every FDB entry learned from any PW is cleared. Using the keywords `vpls_name` clears every FDB entry, (both PW and front panel Ethernet port for the service VLAN) associated with the specified VPLS and the associated VLAN. If the specified VPLS is not bound to a VLAN, the following error message appears:

```
Error: vpls VPLS_NAME not bound to a vlan
```

Using the keywords `vpls_name` and `peer_ip_address` clears all FDB entries from the PW associated with the specified VPLS and LDP peer.

Once the information is cleared from the FDB, any packet destined to a MAC address that has been flushed from the hardware will be flooded until the MAC address has been re-learned.

Example

This example clears the FDB information for VPLS 1:

```
clear fdb vpls vpls1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure fdb vpls agingtime

```
configure fdb vpls agingtime <seconds>
```

Description

Configures the FDB aging time for MAC addresses learned on pseudo wires.

Syntax Description

<seconds>	Specifies the FDB aging time, in seconds. The minimum is 15 and the maximum is 1 million. To disable the aging timer, configure 0 seconds.

Default

300 seconds (or 5 minutes)

Usage Guidelines

The default VPLS FDB agingtime is the same as the Ethernet FDB agingtime, 300 seconds (or 5 minutes). The minimum VPLS FDB agingtime is 15 seconds and the maximum is 1 million seconds. Setting the VPLS FDB agingtime to zero disables VPLS MAC aging. The FDB agingtime is applied to all VPLS instances.

Example

The following example sets the FDB agingtime to 500 seconds:

```
configure fdb vpls agingtime 500
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure iproute add

```
configure iproute add [<ipaddress> <netmask> | <ipNetmask>] [{<gateway>
{<metric>} {vr <vrname>} {unicast-only | multicast-only}} | {lsp
<lsp_name> {<metric>}}]
```

Description

Assigns a named LSP to a specific IP route.

Syntax Description

<ipaddress>	Specified an IP address.
<netmask>	Specifies an IP address/prefix length.
<ipNetmask>	Specifies an IP address/prefix length.
<gateway>	Specified a gateway
<metric>	Specifies a cost metric.
vr <vrname>	Specifies the virtual router to which the route is added.
unicast-only	Specifies only unicast traffic for the route
multicast-only	Specifies only multicast traffic for the route.
lsp <lsp_name>	Specifies a named MPLS LSP to be used to reach the route.
<metric>	Specifies a cost metric.

Default

N/A.

Usage Guidelines

This command assigns a named LSP to a specific IP route. Once configured, all IP traffic matching the configured route is forwarded over the specified LSP. For an RSVP-TE LSP, the correct label information is only associated with the route if the LSP is active. If the RSVP-TE LSP is disabled or is withdrawn, the label information is removed from the route table and the route entry is marked down. If multiple LSPs are added to a route and ECMP is enabled using route-sharing command, only one LSP is used to forward IP traffic.



NOTE

MPLS LSPs are only supported in VR-Default.

Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure iproute add default

```
configure iproute add default [{<gateway> {<metric>} {vr <vrname>}}
{unicast-only | multicast-only}} | {lsp <lsp_name> {<metric>}}]
```

Description

Assigns a named LSP to the default route.

Syntax Description

<gateway>	Specifies a gateway.
<metric>	Specifies a cost metric.
vr <vrname>	Specifies the virtual router to which the route is added.
unicast-only	Specifies only unicast traffic for the route.
multicast-only	Specifies only multicast traffic for the route.
lsp <lsp_name>	Specifies a named MPLS LSP to be used to reach the default route.
<metric>	Specifies a cost metric.

Default

N/A.

Usage Guidelines

This command assigns a named LSP to the default route. Once configured, all IP traffic matching the configured route is forwarded over the specified LSP. For an RSVP-TE LSP, the correct label information is only associated with the default route if the LSP is active. If the RSVP-TE LSP is disabled or is inactive, the label information is removed from the route table and the default route entry is marked down. If multiple LSPs are added to the default route and ECMP is enabled using the `route-sharing` command, traffic is forwarded over only one LSP.



NOTE

MPLS LSPs are only supported in VR-Default.

Example

The following command specifies a default route for the switch:

```
configure iproute add default lsp lsp598
```


History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure iproute delete

```
configure iproute delete [<ipaddress> <netmask> | <ipNetmask>] [{<gateway>
{vr <vrname>}} | lsp <lsp_name>]
```

Description

Deletes a route entry from the routing table based on the configured LSP.

Syntax Description

<ipaddress>	Specified an IP address.
<netmask>	Specifies an IP address/prefix length.
<ipNetmask>	Specifies an IP address/prefix length.
<gateway>	Specifies a gateway.
vr <vrname>	Specifies the virtual router from which the route is deleted.
lsp <lsp_name>	Specifies the named MPLS LSP used to reach the route.

Default

N/A.

Usage Guidelines

This command deletes a route entry from the routing table based on the configured LSP. If the configured IP netmask and LSP name do not match any route entry, the command fails and nothing is deleted.



NOTE

MPLS LSPs are only supported in VR-Default.

Example

The following command deletes an address from the LSP:

```
configure iproute delete 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure iproute delete default

```
configure iproute delete default [{<gateway> {vr <vrname>}} | lsp
<lsp_name>]
```

Description

Deletes a default route entry from the routing table based on the configured LSP.

Syntax Description

<gateway>	Specifies a gateway.
vr <vrname>	Specifies the virtual router from which the route is deleted
lsp <lsp_name>	Specifies the named MPLS LSP used to reach the default route.

Default

N/A.

Usage Guidelines

This command deletes a default route entry from the routing table based on the configured LSP. If the specified LSP name doesn't match any configured default route entry, the command fails and nothing is deleted.



NOTE

MPLS LSPs are only supported in VR-Default.

Example

The following command deletes a default route from the LSP:

```
configure iproute delete default lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls add vlan

```
configure mpls add [{vlan} <vlan_name> | vlan all]
```

Description

Adds an MPLS interface to the specified VLAN.

Syntax Description

vlan <vlan_name>	Identifies the VLAN where the MPLS interface will be added.
vlan all	Adds the MPLS interface to all VLANs that have a configured IP address.

Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. By default, MPLS, LDP, and RSVP-TE are disabled for the MPLS interface. The specified VLAN should have an IP address configured and should have IP forwarding enabled. The MPLS interface on the VLAN will not become active until these two conditions are met. Also, if the IP address is unconfigured from the VLAN or IP forwarding is disabled for the VLAN, the MPLS interface will go down. The MPLS interface state is viewed using the `show mpls interface` command. When the *all VLANs* option is selected, an MPLS interface is added to all VLANs that are currently configured with an IP address and IP forwarding enabled.

The VLAN must be operational for the MPLS interface to be “up.” This means that at least one port in the VLAN must be active or the VLAN must be enabled for loopback mode.

It is recommended that when you configure MPLS on an OSPF interface that can be used to reach a given destination, you should configure MPLS on all OSPF interfaces that can be used to reach that destination. (You should enable MPLS on all of the VLANs connected to the backbone network).

Example

The following example adds MPLS to the VLAN “vlan-usa.”

```
configure mpls add vlan vlan_usa
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls delete vlan

```
configure mpls delete [{vlan} <vlan_name> | vlan all]
```

Description

Removes an MPLS interface from the specified VLAN.

Syntax Description

vlan <vlan_name>	Identifies the VLAN where the MPLS interface will be deleted.
vlan all	Deletes the MPLS interface from all VLANs that have MPLS configured.

Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. If the MPLS interface is deleted, all configuration information associated with the MPLS interface is lost. Issuing this command brings down all LDP neighbor sessions and all LSPs that are established through the specified VLAN interface. When the *all VLANs* option is selected, the MPLS interface for all MPLS configured VLANs is deleted.

Example

The following example deletes MPLS from the VLAN "vlan_k."

```
configure mpls delete vlan vlan_k
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls exp examination

```
configure mpls exp examination {value} <value> {qosprofile} <qosprofile>
```

Description

Configures the QoS profile that will be used for the EXP value when EXP examination is enabled.

Syntax Description

deny	
value <value>	Specifies the value that will be used for the EXP value.
qosprofile <qosprofile>	Specifies the QoS profile that will be used for the EXP value.

Default

The QoS profile matches the EXP value + 1

Usage Guidelines

This command configures the QoS profile that will be used for the EXP value when EXP examination is enabled. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance. The EXP value must be a valid number from 0 through 7 and the qosprofile must match one of the switch's QoS profiles.



NOTE

EXP examination must be enabled using the “enable mpls exp examination” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets QoS profile q5 to be used for EXP value 7:

```
configure mpls exp examination value 7 qosprofile 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls exp replacement

```
configure mpls exp replacement {qosprofile} <qosprofile> {value} <value>
```

Description

Configures the EXP value that will be used for the specified QoS profile when EXP replacement is enabled.

Syntax Description

deny	
qosprofile <qosprofile>	Specifies the QoS profile that will be used for the EXP value.
value <value>	Specifies the value that will be used for the EXP value.

Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command configures the EXP value that will be used for the QoS profile when EXP replacement is enabled. By default, the EXP value matches the QoS profile - 1. That is, QoS profile qp1 is mapped to EXP value of 0, QoS profile qp2 is mapped to EXP value of 1, etc. This configuration has switch-wide significance. The qosprofile must match one of the switch's QoS profiles and the EXP value must be a valid number from 0 through 7.



NOTE

EXP replacement must be enabled using the “enable mpls exp replacement” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets EXP value 2 to be used with QoS profile 4:

```
configure mpls exp replacement qosprofile qp4 value 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls ldp advertise

```
configure mpls ldp advertise [{direct [all | lsr-id | none]} | {rip [all | none]} | {static [all | none]}]
```

Description

Configures a filter to be used by LDP when originating unsolicited label mapping advertisements to LDP neighbors.

Syntax Description

direct	Specifies that the advertisement filter is applied to the associated FECs with directly-attached routing interfaces.
rip	Specifies that the advertisement filter is applied to FECs associated with RIP routes exported by OSPF.
static	Specifies that the advertisement filter is applied to FECs associated with static routes.
all	Specifies that unsolicited label mapping advertisements are originated for all routes of the specified type.
lsr-id	Specifies that an unsolicited label advertisement is originated for a direct route that matches the MPLS LSR ID.
none	Specifies that no unsolicited label mapping advertisements are originated for the specified route type.

Default

None—the default setting for RIP and static routing methods.

All—the default setting for direct routes.

Usage Guidelines

You can configure how the advertisement filter is applied, as follows:

- **direct**—The advertisement filter is applied to the FECs associated with directly-attached routing interfaces.
- **rip**—The advertisement filter is applied to the FECs associated with RIP routes exported by OSPF.
- **static**—The advertisement filter is applied to the FECs associated with static routes.

You can configure the advertisement filter, as follows:

- **all**—Label mappings are originated for all routes of the specified type. This is the default setting for direct routes.
- **none**—No label mappings are originated for all routes of the specified type. This is the default setting for RIP and static routes.
- **lsr-id**—A label mapping is originated for a direct route that matches the MPLS LSR ID.

Advertising labels for a large number of routes may increase the required number of labels that must be allocated by LSRs. Take care to ensure that the number of labels advertised by LERs does not overwhelm the label capacity of the LSRs.

Example

The following command configures LDP to originate labels for all local IP interfaces:

```
configure mpls ldp advertise direct all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls ldp loop-detection

```
configure mpls ldp loop-detection [{hop-count <hop_count_limit>}  
{path-vector <path_vector_limit>}]
```

Description

Configures the loop-detection parameters used by LDP.

Syntax Description

hop-count	Configures the number of LSRs that the label message can traverse.
hop_count_limit	Specifies the hop count limit. The valid configuration range is from 1 to 255.
path-vector	Configures the maximum number of LSR IDs that can be propagated in the label message.
path_vector_limit	Specifies the path vector limit. The valid configuration range is from 1 to 255.

Default

The default for “hop-count” and “path-vector” limits is 255.

Usage Guidelines

Configuration changes are only applicable to newly created LDP sessions. Disabling and enabling LDP will force all the LDP sessions to be recreated. LDP loop detection must first be enabled for these configuration values to be used.

Example

This command sets the LDP hop count loop detection value to 10. The configured path vector value remains at 255.

```
configure mpls ldp loop-detection hop-count 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls ldp timers

```
configure mpls ldp timers [targeted | link] [{hello-time
<hello_hold_seconds>} {keep-alive-time <keep_alive_hold_seconds>}]
```

Description

Configures LDP peer session timers for the switch.

Syntax Description

targeted	Specifies targeted LDP sessions.
link	Specifies link LDP sessions.
hello-time <hello_hold_seconds>	The amount of time (in seconds) that a hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified hello_hold_seconds, then the hello-adjacency is not maintained with that neighboring LSR. The range is 6 to 65,534 seconds.
keep-alive-time <keep_alive_hold_seconds>	The time (in seconds) during which an LDP message must be received for the LDP session with a particular peer LSR to be maintained. If an LDP PDU is not received within the specified session keep_alive_hold_seconds, the corresponding LDP session is torn down. The range is 6 to 65,534 seconds.

Default

link <hello_hold_seconds> – 15 seconds

targeted <hello_hold_seconds> – 45 seconds

link <keep_alive_hold_seconds> – 40 seconds

targeted <keep_alive_hold_seconds> – 60 seconds

Usage Guidelines

The LDP peer hello-adjacency timers are separately configurable for link and targeted LDP sessions. The hello timer parameter specifies the amount of time (in seconds) that a Hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified hello_hold_seconds, then the hello-adjacency is not maintained with that neighboring LSR.

The session keep_alive_hold_seconds parameter specifies the time (in seconds) during which an LDP message must be received for the LDP session to be maintained. The rate at which Keep Alive messages are sent, provided there are no LDP messages transmitted, is 1/6 the configured keep-alive-time. If an LDP PDU is not received within the specified session keep_alive_hold_seconds interval, the corresponding LDP session is torn down. The minimum and maximum values for hold timers are 6 and 65,534, respectively.

Changes to targeted timers only affect newly created targeted sessions. Disabling and then enabling VPLS or LDP will cause all current targeted sessions to be re-created. The default values for the various times are as follows: link <hello_hold_seconds> (15), link <keep_alive_hold_seconds> (40), targeted <hello_hold_seconds> (45), and targeted <keep_alive_hold_seconds> (60). Changes to the link keep-alive timers do not take effect until the LDP session is cycled.

Example

The following command configures link-level LDP hello adjacency hold time to 30 seconds and the keep alive time to 10 seconds:

```
configure mpls ldp timers link hello-time 30 keep-alive-time 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls lsr-id

```
configure mpls lsr-id <ipaddress>
```

Description

Configures the MPLS LSR ID for the switch.

Syntax Description

<code><ipaddress></code>	Specifies an IP address to identify the MPLS LSR for the switch. The MPLS LSR-ID should be configured to the same IP address as the OSPF Router ID.

Default

No LSR ID is configured by default.

Usage Guidelines

LDP, RSVP-TE, and L2 VPNs all use the LSR ID. It is normally set to the OSPF Router ID.

The LSR ID must be configured before MPLS can be enabled. The LSR ID cannot be changed while MPLS is enabled. It is highly recommended that an IP address be configured on a OSPF enabled loopback VLAN that matches the configured LSR ID and OSPF ID. If an LSR ID loopback IP address is configured, OSPF will automatically advertise the LSR ID as a routable destination for setting up LSPs. The LSR ID will remain active if an interface goes down if the LSR-ID is configured as an IP address on a loopback VLAN, as recommended. This significantly enhances network stability and operation of an MPLS network.

Example

The following command configures the LSR ID to 192.168.50.5

```
configure mpls lsr-id 192.168.50.5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te bandwidth committed-rate

```
configure mpls rsvp-te bandwidth committed-rate <committed_bps> [Kbps |
Mbps | Gbps] [{vlan} <vlan_name> | vlan all] {receive | transmit | both}
```

Description

Specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by RSVP-TE LSP reservations.

Syntax Description

<committed_bps>	Specifies a bitrate for the bandwidth to be reserved.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
vlan	Identifies the VLAN where the bandwidth will be reserved.
<vlan_name>	Specifies VLAN where the bandwidth will be reserved.
vlan all	Specifies that the bandwidth will be reserved for all VLANs that have MPLS configured.
receive	Specifies that the bandwidth will be reserved for ingress traffic only.
transmit	Specifies that the bandwidth will be reserved for egress traffic only.
both	Specifies that the bandwidth will be reserved for both ingress and egress traffic.

Default

The default is zero, which means no RSVP-TE LSP bandwidth reservations will be accepted.

If bandwidth is specified without specifying traffic direction, the default is both directions.

Usage Guidelines

This command specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by dynamic RSVP-TE LSP bandwidth reservations. By sub-allocating reserveable bandwidth for RSVP-TE from the VLAN's available bandwidth, the switch can guarantee that as LSPs are established, a minimum amount of CIR bandwidth will be available for other traffic.

The keyword `both` configures the reserved bandwidth for both ingress and egress LSP CIR reservations and overwrites any previous `receive` or `transmit` settings.

Example

The following command reserves 25 Mbps of CIR bandwidth for all RSVP-TE CIR reservations on the specified VLAN:

```
configure mpls rsvp-te bandwidth committed-rate 25 Mbps vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te lsp add path

```
configure mpls rsvp-te lsp <lsp_name> add path [<path_name> | any] {profile
<profile_name>} {primary | secondary}
```

Description

Adds a configured path to the specified RSVP-TE LSP.

Syntax Description

deny	
<lsp_name>	Specifies a name for the RSVP-TE LSP.
<path_name>	Specifies a name for the path to be used by the RSVP-TE LSP.
any	Configures the specified LSP to use any path.
profile <profile_name>	Specifies a profile to be applied to the specified LSP.
primary	Designates the specified path as the primary path. Only one primary path can be configured for an RSVP-TE LSP.
secondary	Designates the specified path as a secondary path. Up to two secondary paths can be configured for an RSVP-TE LSP.

Default

Unless specified, the profile named “default” will be used. If not specified, the path is designated to be “primary.” If not specified and a primary path has already been added, the path is designated to be “secondary.”

Usage Guidelines

This command adds a configured path to the specified RSVP-TE LSP. The *LSP name* parameter is a character string that is to be used to identify the LSP within the switch and must have first been created. The LSP is not signaled until a path is added to the LSP. Up to three paths can be defined for the LSP: one primary and two secondary. All paths are signaled, but only one path is used to forward traffic at any one time. The switch chooses the local MPLS VLAN interface from which to signal the LSP. To force an LSP to use a specific local MPLS interface, configure the local interface IP address as the first ERO in the associated path. The profile name is optional. If omitted, the default profile is applied to the LSP. The path name may be specified or the LSP can be configured to take any path. Only one path can be configured to take any path through the MPLS network.

An RSVP-TE named LSP can consist of up to one primary and two secondary paths. From a control plane perspective, this means one RSVP-TE LSP can actually have three label switch paths signaled and established. From a data plane perspective, only one of the LSP paths can be used to forward data. The primary path is always used if the LSP path is active.

The LSP defaults to primary when not specified. The designation as either primary or secondary is used to determine which LSP next hop is communicated to Route Manager. As long as there is a primary LSP to the specified destination, any configured secondary LSP to the same destination is not used to forward traffic. If there are no primary LSPs active to the destination, any one of the secondary LSPs

may be used to forward traffic. Primary and secondary LSPs must use different paths. Configurations are limited to one primary and two secondary paths per LSP.

All LSPs are signaled as soon as they are configured. If the primary LSP fails, is not configured, or cannot be established after a 5 second delay, one or more of the configured secondary LSPs becomes the "active" LSP to the destination. If at any time a primary LSP is established to the destination, the secondary LSP next hop information is removed from the Route Manager and the secondary path returns to "inactive" state. If both the primary and secondary paths for the LSP fail, and there are no other RSVP-TE LSPs active to the destination, an LDP LSP can be used if available.

Example

This command adds the path "sydney-bypass" to the LSP named "aus" as a secondary path:

```
configure mpls rsvp-te lsp aus add path sydney-bypass secondary
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te lsp delete path

```
configure mpls rsvp-te lsp <lsp_name> delete path [<path_name> | any | all]
```

Description

Deletes a path from the specified RSVP-TE LSP.

Syntax Description

deny	
<lsp_name>	Specifies a name for the RSVP-TE LSP.
<path_name>	Specifies a name for the path to be deleted from the RSVP-TE LSP.
any	Configures the specified LSP to use any path.
all	Deletes all added paths from the specified RSVP-TE LSP.

Default

N/A

Usage Guidelines

This command deletes a path from the specified RSVP-TE LSP. All the added paths can be deleted by specifying the *all* keyword. If the active path is deleted, then one of the other configured paths will become the active path for the LSP. If there are no other defined paths, then the LSP is marked down and cannot be used to forward IP or VPN traffic.

Example

The following command deletes the path called "through-knightsbridge" for the LSP "london":

```
configure mpls rsvp-te lsp london delete path through-knightsbridge
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te lsp path use profile

```
configure mpls rsvp-te lsp <lsp_name> path [<path_name> | any] use profile
<profile_name>
```

Description

Changes the profile that the configured LSP path uses.

Syntax Description

deny

<lsp_name>	Specifies the RSVP-TE LSP.
<path_name>	Specifies the configured RSVP-TE LSP path.
profile	Specifies a profile to be applied to the configured LSP path.
<profile_name>	

Default

N/A

Usage Guidelines

This command changes the profile that the configured LSP path uses.



NOTE

Changing the profile while an LSP is active may cause the LSP to be torn down and re-signaled.

Example

The following command configures the switch to apply the LSP profile “gold-class” to the LSP path “sydney-bypass” for the LSP “aus”:

```
configure mpls rsvp-te lsp aus path sydney-bypass use profile gold-class
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te lsp transport

```
configure mpls rsvp-te lsp <lsp_name> transport [ip-traffic [allow | deny]
| vpn-traffic [allow {all | assigned-only} | deny]]
```

Description

Configures the type of traffic that may be transported across a named LSP

Syntax Description

deny	
<lsp_name>	Specifies the RSVP-TE LSP.
ip-traffic	Controls the forwarding of routed IP traffic across the specified LSP.
vpn-traffic	Controls the forwarding of VPN traffic over the LSP.
allow	Allows transport of the specified traffic across the LSP.
deny	Denies transport of the specified traffic across the LSP.
allow	Allows all VPLS VPN traffic to be transported across the LSP.
all	Allows the transmission of all VPN traffic over the LSP.
assigned-only	Limits the transport of VPN traffic to VPLS instances that are explicitly configured to use the specified LSP name

Default

The default behavior is to allow RSVP-TE LSPs to transport all types of traffic without restriction.

Usage Guidelines

This command configures the type of traffic that may be transported across a named LSP. By default, both IP traffic and VPN traffic are set to allow transport for a newly created LSP. The *ip-traffic* keyword is used to allow or deny forwarding of routed IP traffic across the specified LSP. If allowed, the LSP label information is inserted into the routing table and the switch will forward traffic over the LSP that matches the IP route entry to which this LSP is associated. If denied, the LSP label information is removed from the routing table and the switch will not use the LSP to transport IP traffic. The *vpn-traffic* keyword controls the transmission of VPN traffic over the LSP. When denied, the LSP is not used as a transport for Pseudo Wires or other VPN related traffic. These transport configuration options are independent. For example, if *vpn-traffic* is set to *allow* and *ip-traffic* is set to *deny*, then no routed IP traffic will be transported across the LSP, but the LSP may still be used to transport VPN traffic.

The optional *assigned-only* keyword limits the transport of VPN traffic to only those VPLS instances that are explicitly configured to use the specified LSP name.

Example

The following command prevents the switch from using LSP “aus” to forward IP traffic:

```
configure mpls rsvp-te lsp aus transport ip-traffic deny
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te metric

```
configure mpls rsvp-te metric [<value> | use-igp] {vlan} <vlan_name>
```

Description

Configures the TE metric value for the RSVP-TE interface specified by the *vlan_name* argument.

Syntax Description

<value>	Specifies a value for the RSVP-TE metric.
vlan	Specifies that the RSVP-TE metric will be configured on a VLAN.
<vlan_name>	Specifies VLAN where the RSVP-TE metric will be configured.

Default

The associated default IGP metric.

Usage Guidelines

The TE metric can be any unsigned non-zero 32-bit integer. The default value for the RSVP-TE interface is to use the associated default IGP metric. The TE metric is exchanged between OSPF routers and is used in the calculation of the CSPF topology graph.

Example

The following command configures an RSVP-TE metric of 220 on the specified VLAN:

```
configure mpls rsvp-te metric 220 vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te path add ero

```
configure mpls rsvp-te path <path_name> add ero <ipNetmask> [strict |
loose] {order <number>}
```

Description

Adds an IP address to the Explicit Route Object (ERO) for the specified path name.

Syntax Description

<path_name>	Specifies the path to which the IP address is added.
<ipNetmask>	Specifies an IP prefix.
strict	Specifies that the subobject must be topologically adjacent to the previous subobject in the ERO list.
loose	Specifies that the subobject need not be topologically adjacent to the previous subobject in the ERO list.
order <number>	Specifies the LSR path order

Default

The “order” value defaults to 100 if the path has no EROs configured or a value 100 more than the highest order number configured for the path.

Usage Guidelines

This command adds an IP address to the Explicit Route Object (ERO) for the specified path name. The RSVP-TE routed path may be described by a configured sequence of the LSRs and/or subnets that the path traverses. Each defined LSR or subnet represents an ERO subobject. Up to 64 subobjects can be added to each path name. The *ERO* keyword identifies an LSR using an IP prefix, which may represent an LSR's Router ID, loopback address, or direct router interface. Each IP prefix is included in the ERO as an IPv4 subobject.

If the ERO is specified as *strict*, the *strict* subobject must be topologically adjacent¹ to the previous subobject as listed in the ERO. If the ERO is specified as *loose*, the *loose* subobject is not required to be topologically adjacent to the previous subobject as listed in the ERO. If the specified IP prefix matches the OSPF router ID or a configured loopback IP address, the ERO must be configured as *loose*.

The LSR path order is optionally specified using the *order* keyword. The *order number* parameter is an integer value from 1 to 65535. IP prefixes with a lower number are sequenced before IP prefixes with a higher number. Thus, the LSP path will follow the configured path of IP prefixes with a number value from low to high. If the *order* keyword is not specified, the *number* value for the LSR defaults to a value equal to the current highest number value plus 100. If the list of IP prefixes added to the path does not reflect an actual path through the network topology, the path message will be returned with an error from a downstream LSR and the LSP will not be established.

1. "Topologically adjacent" indicates that the router next hop matches either the interface IP address or OSPF router ID of an immediate peer LSR.

The order of a configured subobject cannot be changed. The ERO subobject must be deleted and re-added with a different order. If a subobject is added to or deleted from the ERO while the associated LSP is established, the path is torn down and is re-signaled using the new ERO. Duplicate ERO subobjects are not allowed.

Defining an ERO for the path is optional. If no ERO is configured, the path is signaled along the best available path and the ERO is not included in the path message. When the last subobject in the ERO of the path message is reached and the egress IP node of the path has not been reached, the remaining path to the egress node is signaled along the best available path. If the next subobject in the ERO is loose, the best available path to the next subobject is chosen. Configuring EROs could lead an LSP to take an undesirable path through the network, so care should be taken when specifying EROs.

Example

The following example adds the IP interface address 197.57.30.7/24 as a loose ERO to the path “sydney-bypass”:

```
configure mpls rsvp-te path sydney-bypass add ero 197.57.30.7/24 loose
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te path delete ero

```
configure mpls rsvp-te path <path_name> delete ero [all | <ipNetmask> |  
order <number>]
```

Description

Deletes a subobject from the Explicit Route Object (ERO) for the specified path name.

Syntax Description

<path_name>	Specifies the path from which the ERO is deleted.
all	Specifies that the entire ERO should be deleted from the named path.
<ipNetmask>	Specifies the ERO subobject to be deleted.
order <number>	Specifies the order number of the ERO subobject to be deleted.

Default

N/A.

Usage Guidelines

This command deletes a subobject from the Explicit Route Object (ERO) for the specified path name. The ERO subobject is specified using an IP prefix or order number. If a subobject is deleted from an ERO while the associated LSP is established, the path is torn down and is re-signaled using a new ERO. The *all* keyword may be used to delete the entire ERO from the path name. When there is no configured ERO, the path is no longer required to take an explicit routed path. The path is then signaled along the best available path and no ERO is included in the path message.

Example

The following command deletes all the configured EROs from the path “sydney-bypass”:

```
configure mpls rsvp-te path sydney-bypass delete ero all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te profile

```
configure mpls rsvp-te profile <profile_name> {bandwidth [best-effort |
[{committed-rate <committed_bps> [Kbps | Mbps | Gbps]] {max-burst-size
<burst_size> [Kb | Mb]} {peak-rate <peak_bps> [Kbps | Mbps | Gbps]]}]
{hold-priority <hold_priority>} {mtu [<number> | use-local-interface]}
{record [enabled | disabled]} {setup-priority <setup_priority>}}
```

Description

Configures an RSVP-TE profile with the specified profile name.

Syntax Description

<profile_name>	Specifies the LSP profile.
bandwidth	Specifies bandwidth reservation.
best-effort	Indicates no bandwidth reservation.
committed-rate <committed_bps>	Specifies the committed bandwidth to be reserved across the MPLS network, in bits per second. The range is from 64 Kbps to 10 Gbps.
peak-rate <peak_bps>	Specifies the maximum bandwidth signaled in bits per second. The range is from 64 Kbps to 10 Gbps.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
max-burst-size <burst_size>	Specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The range is from 0 to 1000 Mb.
Kb	Kilobits
Mb	Megabits
hold-priority <hold_priority>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
setup-priority <setup_priority>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
mtu <number>	Specifies the MTU value for the LSP. The range is from 296 to 9216/
use-local-interface	Specifies that the MTU value is inherited from the local egress VLAN interface.
record	Enables hop-by-hop path recording.
enabled	Causes the Record Route Object (RRO) to be inserted into the path message.
disabled	Specifies that no RRO is inserted into the path message.

Default

Bandwidth: Newly-created profiles are configured as best-effort.

Setup-priority: 7 (lowest)

Hold-priority: 0 (highest)

Path recording: disabled

MTU: use-local-interface

Usage Guidelines

A profile is a set of attributes that are applied to the LSP when the LSP is configured using the `configure mpls rsvp-te lsp` command. A *default* profile is provided which cannot be deleted, but may be applied to any TE LSP. The `profile_name` for the *default* profile is `default`. The *default* profile parameter values are initially set to their respective default values. The maximum number of configurable profiles is 1000.

LSPs may signal reserved bandwidth. By default, newly created profiles are configured to not signal bandwidth requirements and thus are classified as best-effort. If bandwidth needs to be reserved across the MPLS network, the bandwidth parameters specify the desired reserved bandwidth for the LSP. The *committed-rate* specifies the mean bandwidth and the *peak-rate* specifies the maximum bandwidth signaled. The peak-rate must be equal to or greater than the committed-rate. If the peak-rate is not specified, traffic will not be clipped above the committed-rate setting. The rates are specified in bps and must be qualified by Kbps, Mbps, or Gbps. The minimum and maximum bandwidth rates are 64 Kbps and 10 Gbps, respectively. The *max-burst-size* specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The minimum burst size is 0 and the maximum burst size is 1000 Mb.

The *setup-priority* and *hold-priority* are optional parameters indicating the LSP priority. During path set up, if the requested bandwidth cannot be reserved through the LSR, the *setup-priority* parameter is compared to the *hold-priority* of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established. Lower numerical values represent higher priorities. The *setup-priority* range is 0 to 7 and the default value is 7 (lowest). The *hold-priority* range is also 0 to 7 and the default value is 0 (highest). If bandwidth is requested for the LSP, the CSPF calculation uses the available bandwidth associated with the CoS as specified by the *hold-priority*.

The bandwidth, *hold-priority*, and *setup-priority* values are signaled in the path message. If the bandwidth setting is changed, all LSPs using this profile are re-signaled. If the bps setting is decreased, a new path message is sent along the LSP indicating the new reservation. If the bps setting is increased, the LSP is torn down and re-signaled using the new bandwidth reservations.

The `record` command is used to enable hop-by-hop path recording. The enabled keyword causes the Record Route Object (RRO) to be inserted into the path message. The RRO is returned in the RESV Message and contains a list of IPv4 subobjects that describe the RSVP-TE path. Path recording by default is disabled. When disabled, no RRO is inserted into the path message.

The *mtu* keyword optionally specifies the MTU value for the LSP. By default, this value is set to use-local-interface. In the default configuration, the MTU value is inherited from the local egress VLAN interface. The minimum MTU value is 296 and the maximum value is 9216. Path MTU information is carried in the Integrated Services or Null Service RSVP objects and is used by RSVP to perform path MTU identification.



NOTE

Changing any of the profile parameters will cause LSPs using the profile to be torn down and re-signaled. There is no guarantee that the re-signaled LSP will be successfully established. Future ExtremeXOS implementations may support the make-before-break LSP concept.

Example

The following command configures the RSVP-TE profile “gold-class” with a committed bandwidth of 100 Mbps and the setup and hold priorities are both set to 0 (highest priority):

```
configure mpls rsvp-te profile gold-class bandwidth committed-rate 100 mbps hold-  
priority 0 setup-priority 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te timers lsp rapid-retry

```
configure mpls rsvp-te timers lsp rapid-retry {decay-rate <percent>}
{delay-interval <milliseconds>} {retry-limit [<number>]}
```

Description

Configures the timers associated with rapidly retrying failed LSPs.

Syntax Description

decay-rate <percent>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
delay-interval <milliseconds>	Specifies the time (in milliseconds) to wait before attempting to re-signal the LSP.
retry-limit	Specifies the maximum allowed attempts to establish an LSP.
<number>	Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.

Default

Delay interval: 500 milliseconds

Decay rate: 50%

Retry limit: 10

Usage Guidelines

This command configures the timers associated with rapidly retrying failed LSPs. If an LSP fails to establish, the switch will attempt to rapidly retry the setup by sending additional path messages based on the rapid-retry timers. The delay-interval timer specifies the time (in milliseconds) to wait before sending another path message. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time.

For example, if the decay-rate is set to 50 percent and the current delay-interval time is 500 milliseconds, a path message will be retransmitted in 750 milliseconds. If the LSP fails to establish on the next attempt, a path message will be retransmitted after a further decayed delay interval of 1125 milliseconds (1.125 seconds). A per-LSP delay-interval time is maintained for each LSP until the LSP is established. This process of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to rapidly retry the LSP are suspended.

When the switch starts the process of re-signaling the LSP based on the standard-retry timers, the LSP's rapid-retry timers return to the initial configuration settings. If the standard-retry delay-interval time is reached before all of the rapid-retry attempts have completed, the standard-retry mechanisms take over.

The default rapid-retry LSP timer parameter values are 500 milliseconds for the delay-interval, 50 percent for the decay-rate, and a retry-limit of 10. The valid range for delay-interval is 10 to 1000 milliseconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 100. A value of 0 indicates that the LSP will not be re-signaled using the rapid-retry timers.

When `summary-refresh` or `bundle-message` is enabled, the `rapid-retry` timer values are used for resending any message that is not acknowledged.

**NOTE**

RSVP-TE must be disabled before these parameters can be modified.

Example

The following command sets the maximum number of rapid retries to five:

```
configure mpls rsvp-te timers lsp rapid-retry retry-limit 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te timers lsp standard-retry

```
configure mpls rsvp-te timers lsp standard-retry {decay-rate <percent>}
{delay-interval <seconds>} {retry-limit [<number> | unlimited]}
```

Description

Configures the timers associated with the establishment of an LSP.

Syntax Description

decay-rate <percent>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
delay-interval <seconds>	Specifies the time (in seconds) to wait before attempting to re-signal the LSP.
retry-limit <number>	Specifies the maximum allowed attempts to establish an LSP. Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.
unlimited	Allows unlimited attempts to establish an LSP.

Default

Delay interval: 30 seconds

Decay rate: 0%

Retry limit: unlimited

Usage Guidelines

This command configures the timers associated with the establishment of an LSP. If an LSP fails to establish, the LSP will be re-signaled based on the configuration of these timers. The delay-interval timer specifies the time (in seconds) to wait before attempting to re-signal the LSP. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time. For example, if the decay-rate is set to 50 percent and the current delay-interval time is 30 seconds, the LSP is re-signaled in 45 seconds. If the LSP failed to establish on the next attempt, the delay interval would be further decayed to 67 seconds.

A per-LSP delay-interval time is maintained for each LSP until the LSP is established. This operation of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to establish the LSP are suspended.

Disabling and enabling the LSP will reset the LSP's delay-interval time and retry-limit to the initial configuration settings and LSP establishment attempts resume. The default LSP timer parameter values are 30 seconds for delay-interval, with a 0 percent decay-rate, and retry-limit of unlimited. The valid range for delay-interval is 1 to 60 seconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 255 or unlimited. A value of 0 indicates that the LSP will not be re-signaled.

**NOTE**

RSVP-TE must be disabled before these parameters can be modified.

Example

The following command allows unlimited retries for establishing MPLS RSVP-TE LSPs:

```
configure mpls rsvp-te timers lsp standard-retry retry-limit unlimited
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure mpls rsvp-te timers session

```
configure mpls rsvp-te timers session [{bundle-message-time
<bundle_message_milliseconds>} {hello-keep-multiplier <hello_keep_number>}
{hello-time <hello_interval_seconds>}{refresh-keep-multiplier
<refresh_keep_number>} {refresh-time <refresh_seconds>}{summary-refresh-
time <summary_refresh_milliseconds>}] [{vlan} <vlan_name> | vlan all]
```

Description

Configures the RSVP-TE protocol parameters for the specified VLAN.

Syntax Description

bundle-message-time <bundle_message_milliseconds>	Specifies the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The valid range is from 50 to 3000 milliseconds.
hello-keep-multiplier <hello_keep_number>	Specifies the number of hello-time intervals that can elapse before an RSVP-TE peer is declared unreachable. The range is from one to 255.
hello-time <hello_interval_seconds>	Specifies the RSVP Hello packet transmission interval. The valid range is from 1 to 60 seconds.
refresh-keep-multiplier <refresh_keep_number>	Specifies a factor to be used in calculating the maximum allowed interval without an RSVP refresh message before an RSVP session will be torn down. The range is from one to 255.
refresh-time <refresh_seconds>	Specifies the interval for sending refresh path messages. The range is from 1 to 600 seconds.
summary-refresh-time <summary_refresh_milliseconds>	Specifies the interval for sending summary refresh messages. The valid range is from 50 (1/20 second) to 10000 (10 seconds).
vlan <vlan_name>	Identifies the VLAN where the protocol parameters will be configured.
vlan all	Specifies a particular VLAN where the protocol parameters will be configured.
	indicates that the protocol configuration parameters will apply to all RSVP-TE enabled VLANs

Default

Bundle-message-time: 1000 milliseconds (1 second)
Hello-keep-multiplier value: three
Hello-time: 3 seconds
Refresh-keep-multiplier value: three
Refresh-time: 30 seconds
Summary-refresh-time: 3000 milliseconds (3 seconds)

Usage Guidelines

This command configures the RSVP-TE protocol parameters for the specified VLAN. The VLAN keyword *all* indicates that the configuration changes will apply to all VLANs that have been added to MPLS.

The *hello-time* value specifies the RSVP hello packet transmission interval. The RSVP hello packet enables the switch to detect when an RSVP-TE peer is no longer reachable. If an RSVP hello packet is not received from a peer within the configured interval, the peer is declared down and all RSVP sessions to and from that peer are torn down. The formula for calculating the maximum allowed interval is: $[\text{hello-time} * \text{hello-keep-multiplier}]$. The default hello-interval time is 3 seconds with a valid range from 1 to 60 seconds. The default hello-keep-multiplier value is three with a range from one to 255.

The refresh-time specifies the interval for sending refresh path messages. RSVP refresh messages provide "soft state" link-level keep-alive information for previously established paths and enable the switch to detect when an LSP is no longer active. Path messages are used to refresh the LSP if summary refresh is disabled. If summary refresh is enabled, summary refresh messages are sent in place of sending individual path messages for every LSP. The default refresh-time is 30 seconds. The minimum and maximum refresh-time values are one and 600 (or 10 minutes) respectively.

If summary refresh is enabled, summary refresh messages are sent at intervals represented by the configured summary-refresh-time. The configurable summary-refresh-time range is 50 milliseconds (one twentieth of a second) to 10000 milliseconds (10 seconds). The default setting for summary-refresh-time is 3000 milliseconds (3 seconds). RSVP sessions are torn down if an RSVP refresh message is not received from a peer within the configured interval. The formula for calculating the maximum allowed interval is: $[(\text{refresh-keep-multiplier} + 0.5) * 1.5 * (\text{refresh-time or summary-refresh-time})]$. The default refresh-keep-multiplier value is three. The minimum and maximum refresh-keep-multiplier values are one and 255 respectively.

The bundle-message-time, specified in milliseconds, indicates the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The default bundle-message-time is 1000 milliseconds (one second). The bundle-message-time value may be set to any value between 50 milliseconds and 3000 milliseconds (or 3 seconds). Message bundling is only attempted when it is enabled.



NOTE

Summary refresh must be enabled using the "enable mpls rsvp-te summary-refresh" command for a configured summary-refresh-time to actually be used.

Example

The following command sets the RSVP-TE hello time to 5 seconds on all MPLS interfaces:

```
configure mpls rsvp-te timers session hello-time 5 vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls

```
configure vpls <vpls_name> {dot1q [ethertype <hex_number> | tag [include |
exclude]]} {mtu <number>}
```

Description

Configures VPLS parameters.

Syntax Description

vpls_name	Identifies the VPLS within the switch (character string)
dot1q	Specifies the action the switch will perform with respect to the 802.1Q ethertype or tag.
ethertype	Overwrites the ethertype value for the customer traffic sent across the PW
hex_number	Identifies the ethertype, uses the format of 0xN.
tag	Specifies the action the switch will perform with respect to the 802.1Q tag.
include	Includes the 802.1Q tag when sending packets over the VPLS L2 VPN.
exclude	Strips the 802.1Q tag before sending packets over the VPLS L2 VPN.
mtu	Specifies the MTU value of the VPLS transport payload packet.
number	The size (in bytes) of the MTU value. The configurable MTU range is 1492 through 9216. The default VPLS MTU value is 1500.

Default

dot1q tag - included

ethertype - the configured switch ethertype is used.

number (MTU) - 1500

Usage Guidelines

This command configures the VPLS parameters. Pseudo wires (also referred to as PWs) are point-to-point links used to carry VPN traffic between two devices within the VPLS. Each device must be configured such that packets transmitted between the endpoints are interpreted and forwarded to the local service correctly. The optional `ethertype` keyword may be used to overwrite the Ethertype value for the customer traffic sent across the PW. By default, the configured switch ethertype is used. If configured, the ethertype in the outer 802.1q field of the customer packet is overwritten using the configured ethertype value. The ethertype value is ignored on receipt.

Optionally, the switch can be configured to strip the 802.1q tag before sending packets over the VPLS L2 VPN. This capability may be required to provide interoperability with other vendor products or to emulate port mode services. The default configuration is to include the 802.1q tag.

The `mtu` keyword optionally specifies the MTU value of the VPLS transport payload packet (customer packet). The `mtu` value is exchanged with VPLS-configured peer nodes. All VPLS peer nodes must be configured with the same `mtu` value. If the `mtu` values do not match, Pseudo Wires cannot be established between VPLS peers. The `mtu` values are signaled in both the pseudo wire and generalized FECs so that endpoints can verify that MTU settings are equivalent before establishing the PW. By

default the VPLS mtu is set to 1500. The configurable mtu range is 1492 through 9216. Changing the mtu setting will cause established Pseudo Wires to terminate. VPLS payload packets may be dropped if the VPLS MTU setting is greater than the MPLS MTU setting for the PW interface.

Example

The following commands change the various parameters of a particular VPLS:

```
configure vpls vpls1 dot1q ethertype 0x8508
configure vpls vpls1 dot1q ethertype 0x8509 mtu 2500
configure vpls vpls1 dot1q tag exclude mtu 2430
configure vpls vpls1 dot1q mtu 2500
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls add peer

```
configure vpls <vpls_name> add peer <ipaddress>
```

Description

Configures a VPLS peer for the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>ipaddress</code>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP

Default

N/A

Usage Guidelines

Up to 32 peers may be configured for each VPLS. Peers may belong to multiple VPLS instances. The `ipaddress` parameter identifies the peer node that is the endpoint of the VPLS Pseudo Wire. VPLS instances must be configured in a full-mesh. Thus, all peers in the VPLS must have a configured PW to every other peer serving this VPLS. By default, the best LSP will be chosen for the Pseudo Wire. The underlying LSP used by the Pseudo Wire can be configured by specifying the named LSP using the CLI command `configure vpls <vpls_name> peer <ipaddress> [add | delete] mpls lsp <lsp_name>`.

Example

The following command will add a connection from this switch to another switch at 1.1.1.202:

```
configure vpls vpls1 add peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on the BlackDiamond switch only

configure vpls add service

```
configure vpls <vpls_name> add service [{vlan} <vlan_name> | {vman}
<vman_name>]
```

Description

Configures service for VPLS.

Syntax Description

vpls_name	Identifies the VPLS interface within the switch (character string)
vlan_name	Logically binds the VLAN to the specified VPLS.
vman_name	Adds the named VMAN to the VPLS.

Default

N/A

Usage Guidelines

This command configures the VPLS service for the specified `vpls_name`. The VPLS service may be a customer VLAN or a customer VMAN. Specifying the `vlan_name` logically binds the VLAN to the specified VPLS. Only one VLAN or VMAN may be configured per VPLS.

When a VLAN service has been configured for a VPLS, the VLAN is added to the VPLS specified by the `vpls_name`. The VLAN ID is locally significant to the switch. Thus, each VLAN VPLS interface within the VPLS network may have a different VLAN ID service bound to the VPLS. This greatly simplifies VLAN ID coordination between metro network access points. Traffic may be switched locally between VLAN ports if more than one port is configured for the VLAN.

When a VMAN service has been configured for a VPLS, the VMAN is added to the VPLS specified by `vpls_name`. The VMAN ID is locally significant to the switch. Thus, each VMAN VPLS interface within the VPLS network may have a different VMAN ID, just like the VLAN service. The only difference is that the VPLS network overwrites the outer VMAN tag on VPLS egress and leaves the inner VLAN tag unmodified. Because the inner VLAN tag is considered part of the customer packet data, the VMAN service can be used to emulate port-based services. This is accomplished by configuring the VPLS to strip the 802.1Q tag from the tunneled packet. Since the switch inserts the VMAN tag when the packet is received and the 802.1Q tag is stripped before the packet is sent on the VPLS pseudo wire, all packets received on ports that are members of the VMAN are transmitted unmodified across the VPLS. The command `configure vpls <vpls_name> dot1q tag exclude` is used to configure the switch to strip the 802.1Q tag on the VPLS.

Example

The example below adds a VLAN and a VMAN to the named VPLS:

```
configure vpls myvpls add service vlan myvlan
configure vpls myvpls add service vman myvman
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls delete peer

```
configure vpls <vpls_name> delete peer [<ipaddress> | all]
```

Description

Deletes a VPLS peer from the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>ipaddress</code>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP
<code>all</code>	Deletes all VPLS peers.

Default

N/A.

Usage Guidelines

This command deletes a VPLS peer from the specified `vpls_name`. When the VPLS peer is deleted, VPN connectivity to the VPLS peer is terminated. The `all` keyword may be used to delete all peers associated with the specified VPLS.

Example

The following example removes connectivity to 1.1.1.202 from VPLS1:

```
configure vpls vpls1 delete peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls delete service

```
configure vpls <vpls_name> delete service [{vlan} <vlan_name> | {vman}  
<vman_name>]
```

Description

Deletes local VPLS service from the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS interface within the switch (character string)
<code>vlan_name</code>	Logically binds the VLAN to the specified VPLS.
<code>vman_name</code>	Adds the named VMAN to the VPLS.

Default

N/A

Usage Guidelines

This command deletes the local VPLS service from the specified `vpls_name`. Specifying the `vlan_name` or `vman_name` deletes the service from the VPLS. If there are no services configured for the VPLS, all Pseudo Wires within the VPLS are terminated from the switch.

Example

The following example removes a service interface from a VPLS:

```
configure vpls vpls1 delete vman vman1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls peer

```
configure vpls <vpls_name> peer <ipaddress> [limit-learning <number> |
unlimited-learning]
```

Description

Configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer.

Syntax Description

vpls_name	Identifies the VPLS within the switch (character string)
ipaddress	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
limit-learning	Specifies a limit to the number of MAC SAs to be learned for the specified VPLS and peer.
number	The maximum number of MAC SAs that can be learned for the specified VPLS and peer.
unlimited-learning	Specifies no limit to the number of MAC SAs to be learned for the specified VPLS and peer.

Default

Unlimited.

Usage Guidelines

This command configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer. This parameter can only be modified when the specified VPLS is disabled. The `unlimited-learning` keyword can be used to specify that there is no limit. The default value is `unlimited-learning`.

Example

The following example will cause no more than 20 MAC addresses to be learned on VPLS1's pseudo wire to 1.1.1.202:

```
configure vpls vpls1 peer 1.1.1.202 limit-learning 20
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

configure vpls peer mpls lsp

```
configure vpls <vpls_name> peer <ipaddress> [add | delete] mpls lsp
<lsp_name>
```

Description

Configures a named LSP to be used for the PW to the specified VPLS peer.

Syntax Description

vpls_name	Identifies the VPLS within the switch (character string)
ipaddress	Specifies the IP address for the peer node that is the endpoint of the PW-LSP
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the lsp_name parameter from the PW-LSP aggregation list.
lsp_name	Removes the specified lsp.

Default

N/A.

Usage Guidelines

This command configures a named LSP to be used for the PW to the specified VPLS peer. The delete keyword will remove the LSP specified by the lsp_name. If all the named LSPs are deleted to the configured VPLS peer, VPLS will attempt to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the VPLS peer, VPN connectivity to the VPLS peer is lost. Currently, the VPLS pseudo wire uses only one LSP.

Example

The following examples add and remove a named LSP:

```
configure vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

create mpls rsvp-te lsp

```
create mpls rsvp-te lsp <lsp_name> destination <ipaddress>
```

Description

Creates internal resources for an RSVP-TE LSP.

Syntax Description

<lsp_name>	Identifies the LSP within the switch. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<ipaddress>	Specifies the endpoint of the LSP.

Default

N/A.

Usage Guidelines

This command creates internal resources for an RSVP-TE LSP. The LSP name must begin with a letter followed by up to 31 alphanumeric characters. The *ipaddress* specifies the endpoint of the LSP. The LSP is not signaled until a path is specified for the LSP using the `configure mpls rsvp-te lsp <lsp_name> add path` command. When multiple LSPs are configured to the same destination, IP traffic is load-shared across active LSPs that have IP transport enabled. The maximum number of RSVP-TE LSPs that can be created is 1024.



NOTE

The LSP must be created before it can be configured.

Example

The following command creates an RSVP-TE LSP:

```
create mpls rsvp-te lsp lsp598 destination 11.100.100.8
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

create mpls rsvp-te path

```
create mpls rsvp-te path <path_name>
```

Description

Creates an RSVP-TE routed path resource.

Syntax Description

<path_name>	Identifies the path within the switch. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.

Default

N/A.

Usage Guidelines

This command creates an RSVP-TE path resource. The *path_name* parameter is a character string that is used to identify the path within the switch. The *path_name* string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters. The maximum number of configurable paths is 255.



NOTE

The RSVP-TE LSP is not signaled along the path until an LSP is created and then configured with the specified path_name.

Example

The following example creates an RSVP-TE path:

```
create mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

create mpls rsvp-te profile

```
create mpls rsvp-te profile <profile_name>
```

Description

Creates configured RSVP-TE profile with the specified profile name.

Syntax Description

<profile_name>	Identifies the RSVP-TE profile. The character string must begin with an alphabetic character and may contain up to 31 additional alphanumeric characters.

Default

N/A.

Usage Guidelines

This command creates a configured RSVP-TE profile with the specified profile name. The *default* profile cannot be deleted. If a profile is associated with a configured LSP, the profile cannot be deleted.

Example

The following command creates an RSVP-TE profile:

```
create mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

create vpls fec-id-type pseudo-wire

```
create vpls <vpls_name> fec-id-type pseudo-wire <pwid>
```

Description

Creates a VPLS instance with the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string). The <code>vpls_name</code> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<code>pwid</code>	Specifies a Pseudo Wire ID. Must be a non-zero 32-bit value that has network-wide significance.

Default

For the VPLS dot1q tag, the default value is “exclude.”

Usage Guidelines

This command creates a VPLS instance with the specified `vpls_name`. Each VPLS represents a separate virtual switch instance (VSI). The `vpls_name` parameter is a character string that is to be used to identify the VPLS within the switch. The `vpls_name` string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters. Each VPLS is a member of a single VPN and each VPN may have only one associated VPLS per switch. External to the switch, each VPN has an identifier.

Any non-zero 32-bit value that has network-wide significance can be specified for the identifier. This `pwid` will be used on all pseudo-wires in the VPLS.



NOTE

The switch's LSR ID must be configured before a VPLS can be created.

Example

This example creates a VPLS with 99 as the pseudo wire ID:

```
create vpls vpls1 fec-id-type pseudo-wire 99
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

delete mpls rsvp-te lsp

```
delete mpls rsvp-te lsp [<lsp_name> | all]
```

Description

Deletes internal resources for the specified RSVP-TE LSP.

Syntax Description

<lsp_name>	Specifies the LSP within the switch to be deleted.
all	Deletes all RSVP-TE configured LSPs.

Default

N/A.

Usage Guidelines

This command deletes internal resources for the specified RSVP-TE LSP. The LSP is first withdrawn if it is currently active. Deleting an LSP may cause a Pseudo Wire to fail. Any static routes configured to a deleted LSP are also removed.

Example

The following command deletes the configured RSVP-TE LSP named “lsp598”:

```
delete mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

delete mpls rsvp-te path

```
delete mpls rsvp-te path [<path_name> | all]
```

Description

Deletes a configured RSVP-TE routed path with the specified path name.

Syntax Description

<path_name>	Specifies a path within the switch to be deleted.
all	Deletes all paths not associated with an LSP.

Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE routed path with the specified name. All associated configuration information for the specified path is deleted. If the *all* keyword is specified, all paths not associated with an LSP are deleted.



NOTE

A path cannot be deleted as long as the path name is associated with an LSP.

Example

The following command deletes the configured RSVP-TE path named “path 598”:

```
delete mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

delete mpls rsvp-te profile

```
delete mpls rsvp-te profile [<profile_name> | all]
```

Description

Deletes a configured RSVP-TE profile with the specified profile name.

Syntax Description

<profile_name>	Specifies a configured RSVP-TE profile to be deleted.
all	Deletes all profiles not associated with an LSP, except the <i>default</i> profile.

Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE profile with the specified profile name. If the *all* keyword is specified, all profiles not associated with an LSP are deleted (except for the *default* profile).



NOTE

A profile cannot be deleted as long as the profile name is associated with a configured LSP.



NOTE

The default profile cannot be deleted.

Example

The following command deletes the configured RSVP-TE profile named “prof 598”:

```
delete mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

delete vpls

```
delete vpls [<vpls_name> | all]
```

Description

Deletes the VPLS with the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

N/A.

Usage Guidelines

This command deletes the VPLS with the specified `vpls_name`. All VC-LSPs established to VPLS peers are terminated. The `all` keyword may be used to indicate that all VPLS instances are to be deleted.

Example

This commands deletes the VPLS “myvpls:”

```
delete vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable bgp mpls-next-hop

```
disable bgp mpls-next-hop
```

Description

Disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. ("Calculated" refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.

Example

The following command disables BGP's use of MPLS LSPs to reach BGP routes:

```
disable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable iproute mpls-next-hop

```
disable iproute mpls-next-hop
```

Description

Disables IP forwarding over MPLS LSPs for the default VR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over MPLS LSPs for the default VR. When disabled, any route with an MPLS LSP as its next hop will be inactive and will not be used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.

Example

This command disables IP forwarding over MPLS LSPs.

```
disable iproute mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls

```
disable mpls
```

Description

Disables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When MPLS is disabled, no label traffic is received or transmitted, and all MPLS-related protocol peer sessions are terminated.

Example

The following command globally disables MPLS on the switch:

```
disable mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls exp examination

```
disable mpls exp examination
```

Description

Disables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

When disabled, all received MPLS packets are assigned to QoS profile `qp1`.

Example

The following command disables the assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
disable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls exp replacement

```
disable mpls exp replacement
```

Description

Disables setting an MPLS packet's EXP value based on the packet's QoS profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the `configure mpls exp replacement` command.

When disabled, all MPLS packets will be transmitted with an EXP value of zero.

Example

The following command disables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
disable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls ldp

```
disable mpls ldp [{vlan} <vlan_name> | vlan all]
```

Description

Disables LDP for the specified MPLS-configured VLANs.

Syntax Description

Default

vlan	Disables LDP for one or more specific VLANs.
<vlan_name>	Disables LDP on the specified VLAN.
vlan all	Disables LDP for all VLANs that have been added to MPLS.

Disabled.

Usage Guidelines

When LDP is disabled, all LDP-advertised labels are withdrawn and all LDP peer sessions are terminated on the specified VLAN(s). By default, LDP is disabled for all VLANs. Specifying the optional *all* keyword disables LDP for all VLANs that have been added to MPLS.

Example

The following command disables LDP for all VLANs:

```
disable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls ldp bgp-routes

```
disable mpls ldp bgp-routes
```

Description

Disables LDP's use of IP prefixes learned from BGP when establishing LDP LSPs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables LDP's establishment of LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that MPLS LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

When enabled, LDP uses routes learned via BGP when establishing LDP LSPs. As each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

Example

The following command disables the use of BGP routes by LDP:

```
disable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls ldp loop-detection

```
disable mpls ldp loop-detection
```

Description

Disables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge-capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

When LDP loop detection is disabled, LDP will not attempt to detect routing loops.

Example

The following command globally disables LDP loop detection on the switch:

```
disable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls php

```
disable mpls php [{vlan} <vlan_name> | vlan all]
```

Description

Disables penultimate hop popping (PHP) on the specified VLAN. When enabled, PHP is requested on all LSPs advertised over that VLAN for which the switch is the egress LSR.

Syntax Description

vlan	Disables PHP for one or more specific VLANs.
<vlan_name>	Disables PHP on the specified VLAN.
vlan all	Disables PHP for all VLANs that have been added to MPLS.

Default

Disabled

Usage Guidelines

When PHP is disabled on a VLAN, penultimate hop popping is not requested on any LSPs advertised over that VLAN for which the switch is the egress LSR. Therefore, the Implicit Null Label is not used for any advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the *all VLANs* option is selected, PHP is disabled on all existing MPLS interfaces.



NOTE

PHP is sometimes used to reduce the number of MPLS labels in use. If PHP is enabled on any MPLS interface, a unique MPLS label will be consumed for every label advertised over that interface. Therefore, if PHP is being disabled to reduce label consumption, it should be done on all interfaces for minimal label consumption.

In ExtremeXOS, this command can be executed while MPLS is enabled.

Example

The following command disables penultimate hop popping (PHP) on the specified VLAN:

```
disable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls protocol ldp

```
disable mpls protocol ldp
```

Description

Disables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When LDP is disabled, all advertised LDP labels are withdrawn and LDP peer sessions are terminated. Note that this includes any LDP peer sessions established for L2 VPNs. By default, LDP is globally disabled. While LDP is transitioning to the enabled state, only the MPLS show commands are accepted.

Example

The following command globally disables LDP on the switch:

```
disable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls protocol rsvp-te

```
disable mpls protocol rsvp-te
```

Description

Disables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. While RSVP-TE is transitioning to the disabled state, only the MPLS show commands are accepted.

Example

The following command globally disables RSVP-TE on the switch:

```
disable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls rsvp-te

```
disable mpls rsvp-te te [{vlan} <vlan_name> | vlan all]
```

Description

Disables RSVP-TE for the specified MPLS-configured VLAN.

Syntax Description

vlan	Specifies that RSVP-TE will be disabled on one or more VLANs.
<vlan_name>	Specifies one VLAN where RSVP-TE will be disabled.
vlan all	Disables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables RSVP-TE for the specified MPLS configured VLANs. When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. By default, RSVP-TE is disabled for all MPLS configured VLANs. Specifying the optional `all` keyword disables RSVP-TE for all VLANs that have been added to MPLS.

Example

The following command disables RSVP-TE on the named VLAN:

```
disable mpls rsvp-te vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls rsvp-te bundle-message

```
disable mpls rsvp-te bundle-message [{vlan} <vlan_name> | vlan all]
```

Description

Disables the bundling of RSVP-TE messages for the specified VLAN interface.

Syntax Description

vlan	Specifies that message-bundling will be disabled on one or more VLAN interfaces.
<vlan_name>	Specifies one VLAN interface where message bundling will be disabled.
vlan all	Specifies that message bundling will be disabled on all VLAN interfaces that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables the bundling of RSVP-TE messages for the VLAN specified interface. By default, message bundling is disabled. Specifying the `all` keyword will disable message bundling on all VLANs that have been added to MPLS.

Example

The following command disables message bundling on the specified VLAN:

```
disable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls rsvp-te lsp

```
disable mpls rsvp-te lsp [<lsp_name> | all]
```

Description

Disables an RSVP-TE LSP.

Syntax Description

<lsp_name>	Specifies the LSP within the switch to be disabled.
all	Disables all RSVP-TE configured LSPs.

Default

Enabled.

Usage Guidelines

This command disables an RSVP-TE LSP. When an RSVP-TE LSP is disabled, the switch terminates the LSP by signaling the destination by sending a PATH_TEAR message. If there are other LSPs configured to the same destination, traffic may continue to be transmitted to the destination over another LSP. Disabling an LSP does not otherwise change its configuration.

Example

The following command disables the LSP named “lsp598”:

```
disable mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls rsvp-te summary-refresh

```
disable mpls rsvp-te summary-refresh [{vlan} <vlan_name> | vlan all]
```

Description

Disables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified VLAN interface.

Syntax Description

vlan	Specifies that summary refresh messages cannot refresh the RSVP-TE path state on one or more VLAN interfaces.
<vlan_name>	Specifies one VLAN interface where RSVP-TE summary refresh messages will be disabled.
vlan all	Specifies that summary refresh messages will be disabled on all VLAN interfaces that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables the sending of summary refresh messages to refresh RSVP-TE path state for the specified VLAN interface. By default, summary refresh is disabled. Specifying the `all` keyword will disable summary refresh on all VLANs that have been added to MPLS.

Example

The following command disables summary refresh on the specified VLAN:

```
disable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable mpls vlan

```
disable mpls [{vlan} <vlan_name> | vlan all]
```

Description

Disables the MPLS interface for the specified VLAN(s).

Syntax Description

vlan	Disables an MPLS interface for one or more specific VLANs.
<vlan_name>	Disables an MPLS interface on the specified VLAN.
vlan all	Disables an MPLS interface for all VLANs that have been added to MPLS.

Default

The MPLS interface is disabled for a VLAN.

Usage Guidelines

Disabling MPLS causes all LSPs to be released and all LDP and RSVP-TE peer sessions to be terminated on the specified VLAN(s).

Example

The following command disables an MPLS interface for the specified VLAN:

```
disable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable ospf mpls-next-hop

```
disable ospf mpls-next-hop
```

Description

Disables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF. ("Calculated" refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

Example

The following command disables OSPF's use of MPLS LSPs to reach OSPF routes:

```
disable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable snmp traps mpls

```
disable snmp traps mpls
```

Description

Disables SNMP traps associated with MPLS for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command disables SNMP traps associated with MPLS:

```
disable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable vpls

```
disable vpls [<vpls_name> | all]
```

Description

Disables the VPLS instance specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command disables the VPLS instance specified by `vpls_name`. When a VPLS instance is disabled, all sessions to the configured peers for the VPLS are terminated. If this is an MTU/CE switch, all configured services for this VPLS are immediately isolated from other devices residing on the VPLS VPN. If this is a PE switch, then all MTU/CE switches connected to this PE switch are isolated and unable to communicate with each other or any other MTU/CE switch in the VPLS VPN.

Example

The following command disables the VPLS named “myvpls:”

```
disable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

disable vpls service

```
disable vpls [<vpls_name> | all] service
```

Description

Disables the configured VPLS services for the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS is removed from all peer sessions. The keyword `all` disables services for all VPLS instances.

Example

The following command disables the configured VPLS services for the specified VPLS:

```
disable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable bgp mpls-next-hop

```
enable bgp mpls-next-hop
```

Description

Enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. ("Calculated" refers to an LSP that only reaches part of the way to the destination). IP forwarding over MPLS LSPs must be enabled to forward over calculated LSPs. By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.

Example

The following command enables BGP's use of MPLS LSPs to reach BGP routes:

```
enable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable iproute mpls-next-hop

```
enable iproute mpls-next-hop
```

Description

Enables IP forwarding over MPLS LSPs for the default VR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over MPLS LSPs for the default VR. When enabled, LSP next hops can be used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.

Example

The following command enables IP forwarding over MPLS LSPs:

```
enable iproute mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls

```
enable mpls
```

Description

Enables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Enabling MPLS allows MPLS processing to begin for any enabled MPLS protocols (RSVP-TE and/or LDP).

While MPLS is transitioning to the enabled state, only the MPLS show commands are accepted.

Example

The following command globally enables MPLS on the switch:

```
enable mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls exp examination

```
enable mpls exp examination
```

Description

Enables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value. The EXP values to QoS profile mappings are configured using the `configure mpls exp examination` command.

When disabled, all received MPLS packets are assigned to QoS profile `qp1`.

Example

The following command enables assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
enable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls exp replacement

```
enable mpls exp replacement
```

Description

Enables setting an MPLS packet's EXP value based on the packet's QoS profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the `configure mpls exp replacement` command.

When disabled, all MPLS packets will be transmitted with an EXP value of zero.

Example

The following command enables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
enable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls ldp

```
enable mpls ldp [{vlan} <vlan_name> | vlan all]
```

Description

Enables LDP for the specified MPLS configured VLANs.

Syntax Description

Default

vlan	Enables LDP for one or more specific VLANs.
<vlan_name>	Enables LDP on the specified VLAN.
vlan all	Enables LDP for all VLANs that have been added to MPLS.

Disabled.

Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on the enabled VLAN. Once a peer session is established, LDP will advertise labels for local IP interfaces and for routes learned from other neighboring routers. By default, LDP is disabled for all VLANs that have been added to MPLS. Specifying the optional *all* keyword enables LDP for all MPLS configured VLANs.

Example

The following command enables LDP for all VLANs that have been added to MPLS:

```
enable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls ldp bgp-routes

```
enable mpls ldp bgp-routes
```

Description

Enables LDP to use IP prefixes learned from BGP when establishing LDP LSPs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows LDP to use routes learned via BGP when establishing LDP LSPs. Because each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

When disabled, LDP will not establish LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that MPLS LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

Example

The following command enables the use of BGP routes by LDP:

```
enable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls ldp loop-detection

```
enable mpls ldp loop-detection
```

Description

Enables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

Example

The following command globally enables LDP loop detection on the switch:

```
enable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls php

```
enable mpls php [{vlan} <vlan_name> | vlan all]
```

Description

Enables penultimate hop popping (PHP) on the specified VLAN. When enabled, PHP is requested on all LSPs advertised over that VLAN for which the switch is the egress LSR.

Syntax Description

vlan	Enables PHP for one or more specific VLANs.
<vlan_name>	Enables PHP on the specified VLAN.
vlan all	Enables PHP for all VLANs that have been added to MPLS.

Default

Disabled

Usage Guidelines

Penultimate hop popping is requested by assigning the Implicit Null Label in an advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the *all VLANs* option is selected, PHP is enabled on all configured VLANs that have been added to MPLS.



NOTE

ExtremeWare always used the Implicit NULL Label in conjunction with routes exported into OSPF, regardless of the setting of the PHP configuration parameter. This method of operation is not utilized in ExtremeXOS.

Example

The following command enables penultimate hop popping (PHP) on the specified VLAN:

```
enable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls protocol ldp

```
enable mpls protocol ldp
```

Description

Enables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on VLAN interfaces where LDP has been enabled (see 6.3.5 page 77). Once a peer session is established, LDP can advertise labels for local IP interfaces and for routes learned from other neighboring routers. While LDP is transitioning to the enabled state, only the MPLS show commands are accepted.

Note that the LDP protocol must be enabled to establish VPLS pseudo-wires even if the transport LSPs are being established using RSVP-TE.

Example

The following command globally enables LDP on the switch:

```
enable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls protocol rsvp-te

```
enable mpls protocol rsvp-te
```

Description

Enables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, configured LSPs begin the process of TE LSP establishment and VLAN interfaces that have RSVP-TE enabled begin processing RSVP path/reserve messages. By default, RSVP-TE is disabled. While RSVP-TE is transitioning to the enabled state, only the MPLS show commands are accepted.

Example

The following command globally enables RSVP-TE on the switch:

```
enable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls rsvp-te

```
enable mpls rsvp-te [{vlan} <vlan_name> | vlan all]
```

Description

Enables RSVP-TE for the specified MPLS-configured VLAN.

Syntax Description

vlan	Specifies that RSVP-TE will be enabled on one or more VLANs.
<vlan_name>	Specifies one VLAN where RSVP-TE will be enabled.
vlan all	Enables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, TE LSP establishment for configured LSPs begins and the processing of RSVP path/reserve messages from peer LSRs is permitted. By default, RSVP-TE is disabled for all MPLS-configured VLANs. Specifying the optional `all` keyword enables RSVP-TE for all VLANs that have been added to MPLS.

Example

The following command enables RSVP-TE on the specified VLAN:

```
enable mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls rsvp-te bundle-message

```
enable mpls rsvp-te bundle-message [{vlan} <vlan_name> | vlan all]
```

Description

Enables the bundling of RSVP-TE messages for the specified VLAN interface.

Syntax Description

vlan	Specifies that message-bundling will be enabled on one or more VLAN interfaces.
<vlan_name>	Specifies one VLAN interface where message bundling will be enabled.
vlan all	Specifies that message bundling will be enabled on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

Enabling message bundling can improve control plane scalability by allowing the switch to bundle multiple RSVP-TE messages into a single PDU. Not all devices support bundled messages. If the switch determines that a peer LSR, connected to a specific interface, does not support message bundling, the switch reverts to sending separate PDUs for each message on that interface. By default, message bundling is disabled. Specifying the `all` keyword will enable message bundling on all MPLS-configured VLANs.

Example

The following command enables message bundling on the specified VLAN:

```
enable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls rsvp-te lsp

```
enable mpls rsvp-te lsp [<lsp_name> | all]
```

Description

Enables one or more RSVP-TE LSPs.

Syntax Description

<lsp_name>	Specifies the ingress LSP within the switch to be enabled.
all	Enables all RSVP-TE configured ingress LSPs.

Default

Enabled.

Usage Guidelines

When an RSVP-TE LSP is enabled, the switch attempts to set up the LSP by signaling the destination by sending a path message using the assigned path and profile. By default, all newly created LSPs are enabled and can become active when the LSP has been configured. Note that an LSP must be configured with at least one path before it can be signaled.

Example

The following command enables all RSVP-TE-configured LSPs:

```
enable mpls rsvp-te lsp all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls rsvp-te summary-refresh

```
enable mpls rsvp-te summary-refresh [{vlan} <vlan_name> | vlan all]
```

Description

Enables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified VLAN interface.

Syntax Description

vlan	Specifies that summary refresh messages will refresh the RSVP-TE path state on one or more VLAN interfaces.
<vlan_name>	Specifies one VLAN interface where RSVP-TE summary refresh messages will refresh the RSVP-TE path state.
vlan all	Specifies that summary refresh messages will refresh the RSVP-TE path state on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

Enabling summary refresh can improve control plane scalability by refreshing multiple LSPs in a single message. Not all devices support summary refresh. If the switch determines that a peer LSR, connected to a specific interface, does not support summary refresh, the switch reverts to using path messages on that interface. By default, summary refresh is disabled. Specifying the `all` keyword will enable summary refresh on all MPLS-configured VLANs.

Example

The following command enables summary refresh on the specified VLAN:

```
enable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable mpls vlan

```
enable mpls [{vlan}<vlan_name>|vlan all]
```

Description

Enables the MPLS interface for the specified VLAN.

Syntax Description

vlan	Enables an MPLS interface for one or more specific VLANs.
<vlan_name>	Enables an MPLS interface on the specified VLAN.
vlan all	Enables an MPLS interface for all VLANs that have been added to MPLS.

Default

Disabled.

Example

The following command enables an MPLS interface for the specified VLAN:

```
enable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable ospf mpls-next-hop

```
enable ospf mpls-next-hop
```

Description

Enables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF. ("Calculated" refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

Example

The following command enables OSPF's use of MPLS LSPs to reach OSPF routes:

```
enable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable snmp traps mpls

```
enable snmp traps mpls
```

Description

Enables SNMP traps associated with MPLS for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command enables SNMP traps associated with MPLS:

```
enable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable vpls

```
enable vpls [<vpls_name> | all]
```

Description

Enables the VPLS instance specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command enables the VPLS instance specified by `vpls_name`. By default, all newly created VPLS instances are enabled. When enabled, VPLS attempts to establish sessions between all configured peers. Services must be configured and enabled for sessions to be established successfully.

Example

The following command enables the VPLS instance “myvpls:”

```
enable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

enable vpls service

```
enable vpls [<vpls_name> | all] service
```

Description

Enables the configured VPLS services for the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

This command enables the configured VPLS services for the specified `vpls_name`. When services are disabled, the VPLS is withdrawn from all peer sessions. The keyword `all` enables services for all VPLS instances.

Example

The following command enables the configured VPLS services for the specified VPLS instance:

```
enable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

ping mpls lsp

```
ping mpls lsp [<lsp_name> | any <host> | prefix <ipNetmask>]
{reply-mode [ip | mpls]}
{continuous | count <count>}{dont-fragment} {interval <interval>}
{start-size <start-size> {end-size <end-size>}} {ttl <ttl>}
{{from <from>} {next-hop <hopaddress>}}
```

Description

Sends an MPLS ping packet to a FEC over an LSP.

Syntax Description

isp_name	Specifies the LSP on which to send the MPLS echo request.
any	Allows the echo request to be sent over any available LSP.
<host>	Specifies the FEC using an <i>ipaddress</i> or <i>hostname</i> .
prefix	Specifies a prefix.
<ipNetmask>	Specifies the prefix address.
reply-mode	Specifies the return path for the MPLS echo response.
ip	Specifies the return path for the MPLS echo response. Routes back to the sender via an IPv4 packet with the IP Router Alert option. This is the default mode.
mpls	Specifies the return path for the MPLS echo response. Switches the echo reply back to the sender via an IPv4 packet over an LSP.
continuous	Sends pings continuously until the user intervenes.
count <count>	Determines whether size of the packet will increment by one byte for each new MPLS echo request sent.
dont-fragment	Sets the IP <i>do not fragment</i> bit.
interval <interval>	Specifies the time interval (in seconds) between pings.
start-size	Specifies a number of bytes to include as payload data in the MPLS ping packet.
<start-size>	The number of payload data bytes in the MPLS ping packet. The range is from 1 - 1518 (if jumbo frames are disabled) and from 1 - the configured jumbo packet size (if jumbo frames are enabled). The default is 8 bytes.
end-size	Specifies that the size of the packet will increment by one byte for each new MPLS echo request sent, up to a maximum size.
<end-size>	Specifies the maximum size for the MPLS ping packet.
ttl <ttl>	Sets the time-to-live value in the ping packet
from <from>	Specifies the source IP address of the packet.
next-hop	Routes the MPLS echo request along the LSP that traverses the specified node.
<hopaddress>	Specifies the next-hop address.

Default

Destination IP address for MPLS echo request - random, from the 127 and 128 IP address space

IP TTL - 1 second

TTL value in MPLS echo request - 255

Destination UDP port - 3503

Payload data packet size - 8 bytes

Number of pings sent - 4

Usage Guidelines

This command sends an MPLS ping packet to a FEC over an LSP. The ping command, with `mpls` keyword option, can be used to verify data plane connectivity across an LSP. This is useful, since not all failures can be detected using the MPLS control plane. The `lsp` keyword and `lsp_name` parameter may be used to specify the LSP on which to send the MPLS echo request. The `lsp` keyword along with the any keyword allows the echo request to be sent over any available LSP that terminates at host, specified as an `ipaddress` or `hostname`. If no LSP exists to the specified host, the ping command will fail even though an IP routed path may exist. If the optional next-hop is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This option is useful for specifying a specific LSP when multiple LSPs exist to the specified FEC. For RSVP-TE LSPs, the FEC is implied from the LSP configuration. The TTL value in the MPLS Echo Request is set to 255.

By default, the destination IP address of the MPLS echo request is randomly chosen from the 127/8 IP address space and the IP TTL is set to 1. The destination UDP port is 3503 and the sender chooses the source UDP port.

The optional `start-size` keyword specifies the number of bytes to include as payload data in the MPLS ping packet. If no `start_size` parameter is specified, the size of the payload data is eight bytes. The minimum valid `start_size` value is one. The maximum `start_size` value is variable, depending on the type of MPLS ping packet sent, but the total size of the MPLS ping packet cannot exceed the configured jumbo packet size, if jumbo frames are enabled, or 1518 if jumbo frames are disabled. If the `end_size` keyword is specified, the size of the packet will increment by one byte for each new MPLS echo request sent. The next MPLS echo request is not sent until the MPLS echo response for the previous packet is received. This is useful for detecting interface MTU mismatch configurations between LSRs. The switch ceases sending MPLS echo requests when the specified `end_size` value is reached, the MPLS ping is user interrupted, or an MPLS echo response is not received after four successive retries.

The optional `reply-mode` keyword is used to specify the return path for the MPLS echo response. When the `ip` option is specified, the MPLS echo reply is routed back to the sender via an IPv4 packet with the IP Router Alert option. When the `mpls` option is specified, the MPLS echo reply is switched back to the sender via an IPv4 packet over an LSP. If there is no return LSP available, the MPLS echo reply is not sent. This option is useful for determining if there is a failure along the return LSP. If the `reply-mode` is not specified, the `ip` reply mode is used.

Example

The following example shows a ping command and the resulting display:

```
ping mpls lsp prefix 11.100.100.212/32
Ping(MPLS) : 4 packets, 8 data bytes, interval 1 second(s).
98 bytes from 11.100.100.212: mpls_seq=0 ttl=64 time=6.688 ms
98 bytes from 11.100.100.212: mpls_seq=1 ttl=64 time=6.036 ms
98 bytes from 11.100.100.212: mpls_seq=2 ttl=64 time=6.218 ms
98 bytes from 11.100.100.212: mpls_seq=3 ttl=64 time=6.467 ms
```

```
--- ping statistics ---  
4 packets transmitted, 4 received, 0% loss  
round-trip min/avg/max = 6/6/6/ms
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

restart process mpls

```
restart process mpls {msm <slot>}
```

Description

Restarts the MPLS process when it does not respond to the CLI commands.

Syntax Description

<slot>	Specifies the MSM where process should be terminated and restarted. “A” specifies the MSM installed in slot A, and “B” specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Example

The following command restarts the MPLS process:

```
restart process mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show bandwidth pool

```
show bandwidth pool [ingress | egress | duplex] vlan <vlanname>
```

Description

Displays the configured bandwidth pool settings for the specified VLAN.

Syntax Description

ingress	Displays configured bandwidth pool settings for incoming traffic only.
egress	Displays configured bandwidth pool settings for outgoing traffic only.
duplex	Displays configured bandwidth pool settings for traffic in both directions.
vlan <vlanname>	Displays configured bandwidth pool settings only for the specified VLAN.

Default

N/A.

Usage Guidelines

This command displays the configured bandwidth pool settings for a VLAN. Values displayed include the VLAN, maximum reserveable bandwidth (both ingress and egress), and bandwidth reserved by application and by priority level.

Example

The following command displays bandwidth pool settings and accepted bandwidth reservations for all ports:

```
show bandwidth pool duplex vlan vlan_1
```

```
* BD-10K.1 # show bandwidth pool duplex vlan vlan_1
```

Vlan	(mbps)		Rsvd CIRBW		Cmmt		CIRBW			
	Dir	Phy	BE	Limit	Pools	Total	Avail			
vlan_1	Rx	1000	0	1000	300	300	700			
	Tx	1000	0	1000	500	500	500			

Appl	(mbps)		CIRBW Available in Pool (per priority level)							
	Dir	Pool	0	1	2	3	4	5	6	7
mpls	Rx	300	300	300	290	290	290	290	290	290
	Tx	500	500	500	491	491	491	491	491	491

(Rx)-Receive, (Tx)-Transmit (BE)-Best Effort

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show fdb

```
show fdb [vpls | {vpls} <vpls_name>]
```

Description

Displays FDB information learned for VPLS.

Syntax Description

vpls <	
vpls	Displays every FDB entry learned from PWs.
vpls <vpls_name>	Displays every FDB entry (both PW and front panel Ethernet) associated with the specified VPLS and its associated VLAN.

Default

N/A.

Usage Guidelines

Using this command with the required keyword `vpls` only displays all FDB entries learned from the VPLS pseudo wires (PWs). Using the optional keywords `vpls` and `vpls_name` displays all FDB entries (both PW and front panel Ethernet) associated with the specified VPLS and its associated VLAN. If the specified VPLS is not bound to a VLAN, the following error message appears:

```
Error: vpls VPLS_NAME not bound to a vlan
```

Information displayed includes the MAC address, VPLS name, egress VLAN (if applicable), PW peer, Age, and Port List. Showing the FDB entries for the associated service VLAN/VMAN using the `show fdb vlan` command displays the FDB information for both the local VLAN/VMAN ports and the pseudo wire interfaces.

Example

The following command displays FDB information for the specified VPLS:

```
BD-10000.2 # show fdb vpls
Mac                Vlan          Age  Flags          Port / Virtual Port List
-----
00:80:c8:ca:e4:d9   serv(4092)  0021 d m          gv:1.1.1.46

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
        x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
        Blackhole,
        b - Ingress Blackhole, v - MAC-Based VLAN.

Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
FDB Aging time: 300
FDB VPLS Aging time: 300
```

```
BD-10000.3 # show fdb vpls "gv"
```

Mac	Vlan	Age	Flags	Port / Virtual Port List
00:50:04:d6:e4:92	serv(4092)	0022	d m	3:5
00:80:c8:ca:e4:d9	serv(4092)	0025	d m	gv:1.1.1.46

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
 Blackhole,
 b - Ingress Blackhole, v - MAC-Based VLAN.

Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
 FDB Aging time: 300
 FDB VPLS Aging time: 300

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls

show mpls

Description

Displays MPLS configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The *show mpls* command displays the current values for all the MPLS configuration parameters that apply to the entire switch. The parameters displayed include:

- MPLS and MPLS protocol (RSVP-TE and LDP) status.
- SNMP traps configuration.
- EXP examination/replacement configuration.
- the MPLS LSR ID.
- the list of the VLANs which have been added to MPLS.

Example

The following command displays the MPLS configuration parameters for the switch:

```
show mpls
```

```
* BD-10K.30 # show mpls
```

```
MPLS System
```

```

MPLS Admin Status      : Enabled
MPLS Oper Status       : Enabled
RSVP-TE Admin Status   : Enabled
RSVP-TE Oper Status    : Enabled
LDP Admin Status       : Enabled
LDP Oper Status        : Enabled
SNMP Traps             : Disabled
EXP Examination        : Disabled
EXP Replacement        : Disabled
LSR ID                 : 11.100.100.20
MPLS VLANs            : vlan1
                      : vlan2
                      : vlan3

```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls exp examination

```
show mpls exp examination
```

Description

Displays MPLS EXP value to QoS profile mappings and whether MPLS EXP examination is enabled or disabled.

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS EXP value to QoS profile mappings and the status of MPLS EXP examination (enabled or disabled). These values are set using the `configure mpls exp examination qosprofile` command and can be reset using the `unconfigure mpls exp examination` command.

Example

The following is an example of the output of the `show mpls exp examination` command:

```
* BD-10K.29 # show mpls exp examination
```

```
EXP --> QoS Profile mapping:
  00 --> QP1
  01 --> QP2
  02 --> QP3
  03 --> QP4
  04 --> QP5
  05 --> QP6
  06 --> QP7
  07 --> QP8
EXP Examination is disabled
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls exp replacement

```
show mpls exp replacement
```

Description

Displays the MPLS QoS profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled).

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS QoS profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled). These values are set using the `configure mpls exp replacement qosprofile` command and can be reset using the `unconfigure mpls exp replacement` command.

Example

The following is an example of the output of the `show mpls exp replacement` command:

```
* BD-10K.28 # show mpls exp replacement
```

```
QoS Profile --> EXP mapping:
  QP1 --> 00
  QP2 --> 01
  QP3 --> 02
  QP4 --> 03
  QP5 --> 04
  QP6 --> 05
  QP7 --> 06
  QP8 --> 07
EXP Replacement is disabled
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls interface

```
show mpls interface [{vlan} <vlan_name>] {detail}
```

Description

Displays the MPLS interface information. Information is displayed in tabular format for all VLANs that have been added to MPLS.

Syntax Description

vlan	Specifies to display information for one VLAN.
vlan_name	Specifies the name of the VLAN.
detail	Specifies to display additional status information about the interface.

Default

N/A.

Usage Guidelines

Information displayed includes:

- a list of all VLANs added to MPLS.
- MTU size.
- Local interface IP address.
- Number of RSVP-TE neighbors.
- Number of LDP adjacencies.
- RSVP-TE and LDP uptimes.
- MPLS protocols and capabilities configured on each VLAN.

Specifying the optional detail keyword displays the information in verbose form and also includes the operational state for RSVP-TE and LDP. Specifying a VLAN limits the output to that of the individual VLAN.

Example

The following commands display MPLS interface information and detailed information for VLAN 1:

```
show mpls interface vlan vlan1
```

```
* BD-10K.26 # show mpls interface
```

VLAN Name	Local	MTU	RSVP-TE		LDP		Flags
	IP Address		UpTm	#Nbr	UpTm	#Adj	
-----	-----	----	----	----	----	----	-----
vlan1	11.121.96.20	1500	2h	1	2h	1	MRL-IU
vlan2	11.122.96.20	1500	2h	1	2h	1	MRL-IU

```

vlan3          11.123.96.20      1500    --    0    --    0  MRL-I-

Flags: (M) MPLS Enabled, (R) RSVP-TE Enabled, (L) LDP Enabled,
       (P) PHP Enabled, (I) IP Forwarding Enabled, (U) MPLS Operational

* BD-10K.27 # show mpls interface vlan1 detail

VLAN Name : vlan1
  Local IP Address      : 11.121.96.20
  IP Forwarding         : Enabled
  MPLS I/F MTU         : 1500
  PHP Status           : Disabled
  MPLS Admin Status    : Enabled
  MPLS Oper Status     : Enabled
  RSVP-TE Admin Status : Enabled
    Oper Status       : Enabled
    UpTime            : 0d:2h:41m:40s
    # Neighbors       : 1
  LDP Admin Status     : Enabled
    Oper Status       : Enabled
    UpTime            : 0d:2h:41m:40s
    # Link Adjacencies : 1

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls label

```
show mpls {ldp | rsvp-te} label {summary | <label_num> | [advertised |
received] {<label_num>} | received implicit-null}
```

Description

Displays the label information that is used to forward packets within the MPLS network. If no options are specified, tabular information for all active labels is displayed.

Syntax Description

ldp	Specifies only LDP labels will be displayed.
rsvp-te	Specifies only RSVP-TE labels will be displayed.
label	Specifies which labels will be displayed.
summary	Specifies that an informational summary of the labels will be displayed.
advertised	Specifies only labels advertised to other routers will be displayed.
received	Specifies only labels received from other routers will be displayed.
<label_number>	Specifies only labels that match the specified label number will be displayed.
implicit-null	Specifies only implicit-null labels will be displayed

Default

N/A.

Usage Guidelines

This command displays the label information that is used to forward packets within the MPLS network. If no options are specified, tabular information for all active labels except advertised implicit-null labels is displayed. This information includes the incoming and or outgoing label, next hop LSR, and owning MPLS protocol.

The optional keywords `ldp` and `rsvp-te` limit the output to labels for the specified protocol. The optional keywords `advertised` and `received` limit the output to only labels advertised to other routers or received from other routers, respectively. If the optional `label_number` is specified, only labels that match the specified label number (either ingress or egress) are displayed. The `summary` keyword can be used to display an informational summary of the labels. The `implicit-null` keyword can be used to limit output to only the received implicit-null labels.

Example

The following command displays all the MPLS labels except received implicit_null labels:

```
* BD-12K.1 # show mpls label
```

Advertised Label	Destination Mapping	LSP Flags	Peer Label	NHop Type	NextHop	Name
0x80403	1.1.1.210/32	LE	--	vlan	vlan_1	--
0x80404	9.9.9.210/32	LE	--	vlan	temp	--
0x80405	192.99.1.44/32	LE	--	vlan	loopback	--
0x80406	192.99.1.24/32	LE	--	VPLS	--	vpls_1
0x80400	3.1.101.0/24	LE	--	vlan	gtag101	--
0x80401	3.1.102.0/24	LE	--	vlan	gtag102	--
0x80402	192.87.87.0/24	LE	--	vlan	vlan_2	--
0x80407	192.99.1.44/32	RE	--	loc1	192.99.1.44	vlalsp2

Received Label	Destination Mapping	LSP Flags	NextHop	Name
0x80405	192.99.1.24/32	LI	192.99.1.24	vpls_1
0x80406	192.99.1.24/32	RI	192.87.87.24	vlalsp1
0x00010	12.212.0.0/24	LI	192.87.87.24	--
0x00011	15.212.0.0/24	LI	192.87.87.24	--
0x00012	10.20.30.0/24	LI	192.87.87.24	--
0x00013	11.98.96.0/24	LI	192.87.87.24	--
0x00014	11.100.100.100/32	LI	192.87.87.24	--
0x00015	3.1.97.0/24	LI	192.87.87.24	--
0x0003f	11.179.96.0/24	LI	192.87.87.24	--

(M)Multiple Next Hops (L)LDP (R)RSVP-TE
 (T)Transit LSP (I)Ingress to LSP (E)Egress from LSP
 Press <SPACE> to continue or <Q> to quit:

```
Total number of RSVP-TE labels advertised : 1
Total number of RSVP-TE labels received   : 1
Total number of LDP labels advertised      : 7
Total number of LDP labels received        : 8
```

To see the implicit null labels advertised by this LSP to peers via the LDP protocol, use this command:

```
* BD12K.3 # show mpls ldp label advertised implicit-null
```

Prefix	Advertised Label	LDP Peer	PHP Interface
192.83.83.0/24	IMPNULL	192.99.1.2	vlan_1
		192.99.1.44	vlan_2
192.85.85.0/24	IMPNULL	192.99.1.2	vlan_1
		192.99.1.44	vlan_2
192.86.86.0/24	IMPNULL	192.99.1.2	vlan_1
		192.99.1.44	vlan_2
192.87.87.0/24	IMPNULL	192.99.1.2	vlan_1
		192.99.1.44	vlan_2
192.99.1.24/32	IMPNULL	192.99.1.2	vlan_1
		192.99.1.44	vlan_2

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls ldp

```
show mpls ldp
```

Description

Displays summary configuration and status information for LDP. Global status of LDP, LDP session timer configuration, loop detection, and label advertisement status are included in the display output.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the global status of LDP, LDP session timer configuration, loop detection, label advertisement, and LDP-enabled VLANs.

Example

The following command displays summary configuration and status information for LDP:

```
* BD-10K.25 # show mpls ldp
LDP Admin Status      : Enabled
LDP Oper Status       : Enabled
Protocol Version      : v1*
Label Retention Mode  : Liberal*
Label Distribution Method : Downstream Unsolicited*
LDP BGP LSPs         : Enabled

LDP Loop Detection
  Status              : Disabled
  Hop-Count Limit     : 255
  Path-Vector Limit   : 255

LDP Targeted Timers
  Hello Hold          : 45 seconds
  Keep Alive Hold     : 60 seconds

LDP Link Timers
  Hello Hold          : 15 seconds
  Keep Alive Hold     : 40 seconds
```

```
Label Advertisement
  Direct : Matching LSR-ID Only
  Rip    : None
  Static : None
```

```
LDP VLANs : vlan1
           : vlan2
           : vlan3
```

* Indicates parameters that cannot be modified

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls ldp interface

```
show mpls ldp interface {{vlan} <vlan_name>} {detail | counters}
```

Description

Displays LDP information about MPLS interfaces. Summary information is displayed in tabular format for all VLANs that are configured for MPLS.

Syntax Description

vlan	Displays LDP interface information for one VLAN.
vlan_name	Specifies the name of the VLAN.
detail	Displays LDP interface information including LDP control packet counts.
counters	Displays only the LDP control protocol packet counts.

Default

N/A.

Usage Guidelines

Associated with each VLAN, the following information is displayed:

- VLAN name.
- Negotiated hello hold time.
- Next hello message time.
- Uptime.
- MPLS and LDP state flags.
- Number of hello adjacencies.

If the optional `detail` keyword is specified, the information is shown in verbose form and LDP control packet counts are displayed. If the optional `counters` keyword is specified, only the LDP control protocol packet counts are shown.

Example

The following command displays detailed LDP information for the interface associated with VLAN 1:

```
* BD-10K.22 # show mpls ldp interface vlan1 detail
```

```
VLAN Name : vlan1
  Local IP Address      : 11.121.96.20
  MPLS Admin Status    : Enabled
  MPLS Oper Status     : Enabled
  LDP Admin Status     : Enabled
  LDP Oper Status      : Enabled
```



```

LDP UpTime           : 0d:1h:59m:56s
Current Adjacencies   : 1
Negotiated Hello Hold Time : 15000 ms
Time to Send Next Hello : 4060 ms

```

Counter	Link Adjacencies	Targeted Adjacencies
-----	-----	-----
Shutdown Notifications (Rcvd)	0	0
Shutdown Notifications (Sent)	0	0
Failed Session Attempts (NAKs)	0	0
Hello Errors	0	0
Parameters Advertised Errors	0	0
Max PDU Length Errors	0	0
Label Range Errors	0	0
Bad LDP ID Errors	0	0
Bad PDU Length Errors	0	0
Bad Msg Length Errors	0	0
Bad TLV Length Errors	0	0
Bad TLV Value Errors	0	0
Keep-Alive Timeout Errors	0	0

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls ldp lsp

```
show mpls ldp lsp {prefix <ipNetmask>} {ingress | egress | transit}
{detail}
```

Description

Displays the LSP information associated with LDP that is used to forward packets within the MPLS network. If no options are specified, summary information for all LSPs is displayed.

Syntax Description

prefix	Displays information for a single FEC that matches the prefix.
<ipNetmask>	Designates the FEC for which to display information.
ingress	Displays information for LSPs that originate from the switch into the MPLS network.
egress	Displays information for LSPs that terminate at the switch from the MPLS network.
transit	Displays information for LSPs that traverse the switch.
detail	Display detailed LSP information.

Default

N/A.

Usage Guidelines

If no options are specified, this command displays summary information for all LSPs. This information includes:

- The associated FEC prefix
- Incoming label, or outgoing label, or both
- Local VLAN interface associated with FEC prefix (egress only)
- Next hop LSR (ingress and transit)
- Outgoing VLAN interface (ingress and transit)

Optionally, the LSPs displayed can be further qualified by the keywords `ingress`, `egress`, and `transit`. These keywords qualify the LSPs displayed from the perspective of the switch. Ingress LSPs originate from the switch into the MPLS network. Egress LSPs terminate at the switch from the MPLS network. Transit LSPs traverse the switch. If the optional `prefix` keyword is specified, only the LSP information associated with the FEC that matches the prefix is displayed.

If the `detail` keyword is specified, information is displayed in verbose form and includes received packet and byte counts.

Example

The following command displays information from the LDP label mapping database:

```
* BD-10K.5 # show mpls ldp lsp prefix 10.1.1.1/32
Prefix          Adv Label Peer Label Next Hop      VLAN
10.1.1.1/32     0x00130   0x0165f   11.122.96.8   vlan2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls ldp peer

```
show mpls ldp peer {<ipaddress>} {detail}
```

Description

Displays information about the status of the LDP sessions and hello adjacencies for all LDP peers.

Syntax Description

ipaddress	Display session and hello adjacency information for a single LDP peer.
detail	Display additional detailed information related to the session and adjacencies.

Default

N/A.

Usage Guidelines

Summary information is displayed in tabular format.

Information displayed includes:

- LDP peer IP address
- Session operational state
- Uptime
- Number of hello adjacencies

Specifying the LDP peer's ipaddress displays session and hello adjacency information for a single LDP peer. When the `detail` keyword is specified, additional detailed information related to the session and adjacencies is displayed. This information includes: session information, timers, next-hop addresses, adjacency status and counters.

Example

The following command displays MPLS LDP session information for the LDP entity 11.100.100.30:

```
Mariner3.59 # show mpls ldp peer
```

IP Address	LDP Peer	State	Uptime	Adjacencies
11.100.100.30	11.100.100.55:0	Operational	0d:14h:51m:53s	1
11.100.100.30	14.4.0.99:15	Operational	0d:1h:0m:43s	1
11.100.100.30	14.4.0.99:16	Operational	0d:0h:34m:51s	1

```
Adjacencies:
  Index          : 1          Attempted Sessions : 1
  Targeted Peer  : 11.100.100.210:0 Shutdown Notifications : Sent 0 Rcvd 0
```

```
Mariner3.32 # show mpls ldp peer detail
```

```
Peer: 11.100.100.55      Peer label space: 0 (global)
```

```

    Session State           : Operational
    Session Uptime          : 0d:0h:13m:41s
...
Peer: 14.4.0.99             Peer label space: 15
    Session State           : Operational
    Session Uptime          : 0d:0h:57m:4s
...
Peer: 14.4.0.99             Peer label space: 16
    Session State           : Operational
    Session Uptime          : 0d:0h:31m:12s
...

```

BD-10K.3 # show mpls ldp peer 11.121.96.5 detail

Peer: 11.121.96.5

```

    Session State           : Operational
    Session Uptime          : 0d:0h:4m:1s
    Discontinuity Time      : 34019
    Keep Alive Hold Timer   : 40 (remaining: 34.77)
    Label Distribution Method : DU
    Max PDU Length          : 4096
    Unknown Msg Type Errors : 0
    Unknown TLV Errors      : 0
    Next Hop Addr(s)        : 11.121.96.5

```

Adjacencies:

Index	: 3	Attempted Sessions	: 0
Targeted Peer	: Not Targeted	Shutdown Notifications	: Sent 0 Rcvd 0
Admin Status	: Enabled	No Hello Errors	: 0
Operational Status	: Up	Advertisement Errors	: 0
Label Retention Mode	: Liberal	Max PDU Errors	: 0
Hop Count Limit	: Disabled	Bad LDP Identifier Errors	: 0
Path Vector Limit	: Disabled	Bad PDU Length Errors	: 0
Hello Hold Timer	: 15 (remaining: 14)	Bad TLV Length Errors	: 0
VLAN	: vlan1	Malformed TLV Errors	: 0
Interface address	: 11.121.96.20	Bad Message Length Errors	: 0
		Session Rejected Errors	: 0
		Keep Alive Expired Errors	: 0

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te

```
show mpls rsvp-te
```

Description

Displays displays summary configuration and status information for RSVP-TE.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays summary configuration and status information for RSVP-TE. The parameters displayed include:

- Global status of RSVP-TE.
- Configured standard LSP timer values.
- Configured rapid-retry LSP timer values.
- RSVP-TE VLANs.

Example

The following command shows the summary configuration and status information for RSVP-TE:

```
* BD-10K.16 # show mpls rsvp-te
RSVP-TE Admin Status  : Enabled
RSVP-TE Oper Status   : Enabled

LSP Standard-Retry Timers
  Delay-Interval      : 30 seconds
  Decay-Rate          : 50 %
  Retry-Limit         : unlimited

LSP Rapid-Retry Timers
  Delay-Interval      : 500 milliseconds
  Decay-Rate          : 50 %
  Retry-Limit         : 10

RSVP-TE VLANs        : vlan1
                      : vlan2
                      : vlan3
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te bandwidth

```
show mpls rsvp-te bandwidth [{vlan} <vlan_name>] {detail}
```

Description

Displays the reserved bandwidth for each TE LSP by interface.

Syntax Description

vlan <vlan_name>	Displays the reserved bandwidth for each TE LSP associated with the specified VLAN.
detail	Displays the path information in verbose format.

Default

All TE LSPs for all RSVP-TE-enabled interfaces are shown.

Usage Guidelines

This command displays the reserved bandwidth for each TE LSP by interface. By default, all TE LSPs for all RSVP-TE enabled interfaces are shown.

The optional `vlan` keyword limits the display to only those LSPs that have bandwidth reservations against the specified VLAN. Only committed-rate bandwidth is displayed. Bandwidth is displayed as either *received* or *transmitted* bandwidth with respect to the switch.

LSPs are listed using the configured or signaled LSP name. If the LSP name was not included in the setup control messages (which can only occur when using OEM vendor equipment), the LSP is uniquely identified using a concatenated string that includes the tunnel ID and source IP address. Per VLAN, each LSP is listed in descending priority order. That is, the LSPs listed at the top of each VLAN have the highest bandwidth priority and are less likely to be preempted. Bandwidth priority is determined by the signaled hold-priority and the uptime. The TE LSP with a hold-priority of zero and the highest uptime has the highest bandwidth priority and the TE LSP with a hold-priority of seven and the lowest uptime has the lowest bandwidth priority.

Use the `detail` keyword to display detailed information.

Example

The following command displays bandwidth reservation information for the specified VLAN:

```
show mpls rsvp-te bandwidth vlan vlan_1 detail
```

```
* BD-10K.2 # show mpls rsvp-te bandwidth vlan_1 detail
Vlan          Dir  Pool      CIR (per priority level)
              LSP      0      1      2      3      4      5      6      7
-----
vlan_1        Rx    300
              vlalsp1  -      -      10     -      -      -      -      -
              Available 300    300    290    290    290    290    290    290
              Tx    500
              vlalsp2  -      -      9      -      -      -      -      -
              Available 500    500    491    491    491    491    491    491
```

(Rx)Receive Bandwidth (Tx)Transmit Bandwidth

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te interface

```
show mpls rsvp-te interface {{vlan} <vlan_name>} {detail | counters}
```

Description

Displays RSVP-TE information about MPLS interfaces.

Syntax Description

vlan	Display information for one VLAN.
vlan_name	Specifies the name of the VLAN.
detail	Display RSVP-TE information including the interface up time and LDP control packet counts.
counters	Display only the RSVP-TE control protocol packet counts.

Default

N/A.

Usage Guidelines

This command displays RSVP-TE information about MPLS interfaces. Summary information is displayed in tabular format for all VLANs that are configured for MPLS. The following information is displayed:

- VLAN name.
- Bandwidth reserved.
- TE metric.
- Hello interval time.
- Refresh interval time.
- Summary refresh time.
- Bundle message time.
- Uptime.
- Number of neighbors.
- RSVP-TE state information

When the optional `detail` keyword is specified, additional RSVP-TE information is displayed. This additional information includes:

- RSVP-TE hello keep multiplier.
- RSVP-TE refresh keep multiplier.
- RSVP-TE available bandwidth per priority level.
- RSVP-TE control protocol packet counts.

When the optional `counters` keyword is specified, only the RSVP-TE control protocol packet counts are shown.

Example

The following command displays detailed RSVP-TE information for the interfaces associated with VLAN 1:

```
* BD-10K.15 # show mpls rsvp-te interface vlan1 det
```

```
VLAN Name : vlan1
  Local IP Address       : 11.121.96.20
  MPLS Admin Status      : Enabled
  MPLS Oper Status       : Enabled
  RSVP-TE Admin Status   : Enabled
  RSVP-TE Oper Status    : Enabled
  RSVP-TE Up-Time        : 0d:1h:19m:46s
  # Neighbors            : 1
  Receive CIR            : 50000 Kbps
  Transmit CIR           : 50000 Kbps
  TE Metric              : Use IGP Cost/Metric
  Hello Interval         : 3 seconds
  Refresh Time           : 30 seconds
  Hello Keep Multiplier  : 3
  Refresh Keep Multiplier : 3
  Summary Refresh        : Disabled
  Summary Refresh Time   : 3000 milliseconds
  Bundle Message         : Disabled
  Bundle Message Time    : 1000 milliseconds
```

Dir	Pool	CIR Available (per priority level)							
		0	1	2	3	4	5	6	7

Rx	50	50	50	50	50	50	50	50	50
Tx	50	50	50	50	50	50	50	50	40
(Rx)Receive Bandwidth		(Tx)Transmit Bandwidth							

Message	Sent	Received
-----	-----	-----
PATH	165	165
PATH_TEAR	2	3
PATH_ERR	4	18
RESV	160	145
RESV_TEAR	0	2
RESV_ERR	0	0
RESV_CONFIRM	4	1
SUMMARY_REFRESH	0	0
BUNDLE	0	0
HELLO	42	30

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te lsp

```
show mpls rsvp-te lsp {[[destination | origin] <ipaddress>} {detail}
| summary}
```

Description

Displays the LSP information associated with RSVP-TE that is used to forward packets within the MPLS network.

Syntax Description

destination <ipaddress>	Displays only those LSPs that terminate at the specified IP address.
origin <ipaddress>	Displays only those LSPs that originate at the specified IP address.
detail	Displays the LSP information in verbose format.
summary	Displays only LSP summary statistics.

Default

N/A.

Usage Guidelines

This command displays the LSP information associated with RSVP-TE that is used to forward packets within the MPLS network. If no options are specified, summary information for all RSVP-TE LSPs is displayed. In ExtremeXOS 12.0, summary statistics are added to the display for this command, and the optional `summary` keyword is added to display the statistics only. The information included in the summary output will depend upon which, if any, additional keywords you use. If you specify only the `summary` keyword, the output statistics will include:

- LSP count.
- LSP direction relative to the MPLS network (i.e., ingress, egress, or transit).
- Number of LSPs enabled or disabled (ingress only).
- Number of LSPs with no configured path (ingress only).
- Primary/secondary path state (ingress only).

The optional `destination` keyword limits the display to only those LSPs that terminate at the specified IP address. The optional `origin` keyword limits the display to only those LSPs that originate at the specified IP address. Beginning with Extreme XOS Release 12.0, you can use the `origin` keyword without specifying the LSP type. Specifying the optional `detail` keyword displays the LSP information in verbose format. Additional information displayed includes:

- Incoming and/or outgoing MPLS label.
- Configured profile.
- Cycle count.
- Receive packet and byte counters.

- Recorded routes.
- Number of VPLSs using ingress LSPs.
- Retry counts.
- Information about all configured primary and secondary paths.

Example

Use the following command to display information about the RSVP-TE LSPs with the destination IP address 11.100.100.203:

```
* BD-10K.5 # show mpls rsvp-te lsp destination 11.100.100.203
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp2          any              11.100.100.203  pcvlan2          1d    UESR--IV
jefflsp3          any              11.100.100.203  pcvlan2          1d    UEP---IV
```

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

Summary of RSVP-TE LSPs to destination 11.100.100.203

```
Ingress LSPs (Enabled/Disabled)      : 2 (2/0)
Ingress LSPs with no configured path  : 0
Ingress LSP Paths (Up/Down)          : 3 (2/1)
Transit LSPs                         : 0
Egress LSPs                         : 0
```

Use the following command to display information about all RSVP-TE LSPs:

```
* BD-10K.7 # show mpls rsvp-te lsp
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp0          **no path**      11.100.100.216  -----          0    -E----IV
jefflsp1          jeffpath1        11.100.100.204  n/a              0    --PR--IV
jefflsp2          any              11.100.100.203  pcvlan2          1d    UESR--IV
jefflsp3          any              11.100.100.203  pcvlan2          1d    UEP---IV

Egress LSP Name  Source IP      Destination      Receive I/F      UpTm
-----
jefflsp1         11.100.100.203  11.100.100.48   pcvlan2          1d
```

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

Summary of RSVP-TE LSPs

```
Ingress LSPs (Enabled/Disabled)      : 4 (3/1)
Ingress LSPs with no configured path  : 1
Ingress LSP Paths (Up/Down)          : 5 (2/3)
Transit LSPs                         : 0
Egress LSPs                         : 1
```

Use the following command to display information about a specific LSP:

```
* BD-10K.7 # show mpls rsvp-te lsp jefflsp1
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath1        11.100.100.204   n/a               0  --PR--IV
```

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

```
Summary of Ingress RSVP-TE LSPs named jefflsp1
  Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
  Ingress LSPs with no configured path : 0
  Ingress LSP Paths (Up/Down)          : 2 (0/2)
```

Use the following command to summary information about all LSPs:

```
* BD-10K.7 # show mpls rsvp-te lsp summary
```

Summary of RSVP-TE LSPs

Ingress LSPs (Enabled/Disabled)	: 4 (3/1)
Ingress LSPs with no configured path	: 1
Ingress LSP Paths (Up/Down)	: 5 (2/3)
Transit LSPs	: 0
Egress LSPs	: 1

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te lsp (egress/transit)

```
show mpls rsvp-te lsp [egress | transit] [{<lsp_name>}] [{destination |
origin] <ipaddress>] {detail} | summary}
```

Description

Displays the egress and transit LSP information associated with RSVP-TE that is used to forward packets within the MPLS network.

Syntax Description

egress	Displays only LSPs that terminate at the switch from the MPLS network.
egress <lsp_name>	Displays only the named LSP that terminates at the switch from the MPLS network.
transit	Displays only LSPs that traverse the switch.
transit <lsp_name>	Displays only the named LSP that traverses the switch.
destination <ipaddress>	Displays only those LSPs that terminate at the specified IP address.
origin <ipaddress>	Displays only those LSPs that originate at the specified IP address.
detail	Displays the LSP information in verbose format.
summary	Displays only LSP summary statistics.

Default

N/A.

Usage Guidelines

This command displays the LSP information associated with RSVP-TE that is used to forward packets within the MPLS network. If no options are specified, summary information for all RSVP-TE LSPs is displayed. The information included in the summary output will depend upon which, if any, additional keywords you use. If you specify only the `summary` keyword, the output statistics will include:

- LSP count.
- LSP direction relative to the MPLS network (i.e., ingress, egress, or transit).
- Number of LSPs enabled or disabled (ingress only).
- Number of LSPs with no configured path (ingress only).
- Primary/secondary path state (ingress only).

If the optional LSP name parameter is specified, the LSP information is displayed for the named LSP only. The LSPs displayed can be filtered by the optional keywords `egress` and `transit`. These keywords qualify the LSPs displayed from the perspective of the switch. Egress LSPs identify LSPs that terminate at the switch from the MPLS network. Transit LSPs represent LSPs that traverse the switch.

The optional `destination` keyword limits the display to only those LSPs that terminate at the specified IP address. The optional `origin` keyword limits the display to only those LSPs that originate at the

specified IP address. Specifying the optional `detail` keyword displays the LSP information in verbose format. Additional information displayed includes:

- Incoming and/or outgoing MPLS label.
- Configured profile.
- Cycle count.
- Receive packet and byte counters.
- Recorded routes.
- Number of VPLSs using ingress LSPs.
- Retry counts.
- Information about all configured primary and secondary paths.

Example

Use the following command to display information about egress RSVP-TE LSPs:

```
Mariner5.29 # show mpls rsvp-te lsp egress
```

Egress LSP Name	Source IP	Destination	Receive I/F	UpTm
lsp2	11.100.100.203	11.100.100.55	pcvlan2	2d

Summary of Egress RSVP-TE LSPs

Egress LSPs : 1

Use the following command to display detailed information about transit RSVP-TE LSPs:

```
Mariner5.30 # show mpls rsvp-te lsp transit det
```

Transit LSP Name: lsp4

Tunnel ID : 4	Ext Tunnel ID :
11.100.100.212	
LSP ID : 0	UpTime :
2d:0h:47m:11s	
Source IP : 11.100.100.212	Destination IP :
11.100.100.210	
Rx I/F : 12.151.0.55 - m5-o5	
NextHop : 12.152.0.210	
Tx I/F : 12.152.0.55 - m5-o4	
Rx Packets : 86698201	Adv Label : 0x00010
Rx Bytes : 3507726464	Rcv Label : 0x80400
Record Route : Empty	

Transit LSP Name: lsp4

Tunnel ID : 3	Ext Tunnel ID :
11.100.100.210	
LSP ID : 0	UpTime : 2d:0h:47m:8s
Source IP : 11.100.100.210	Destination IP :
11.100.100.212	
Rx I/F : 12.152.0.55 - m5-o4	
NextHop : 12.151.0.212	
Tx I/F : 12.151.0.55 - m5-o5	

```

Rx Packets      : 86695727          Adv Label       : 0x00011
Rx Bytes       : 3507579254        Rcv Label      : 0x8044b
Record Route   : Empty

```

Use the following command to display information about all RSVP-TE LSPs:

```

* BD-10K.5 # show mpls rsvp-te lsp
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
lsp598           path598          11.100.100.8    vlan1              3m  UEP--OIV

Egress LSP Name Source IP      Destination      Receive I/F      UpTm
-----
test_lsp5        11.100.100.8   11.100.100.20   vlan2              8m

Transit LSP Name Source IP      Destination      Receive I/F      Transmit I/F      UpTm
-----
lsp5a8           11.100.100.5   11.100.100.8    vlan1              vlan2              5m

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
       (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
       (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
       (v) VPN Assigned Traffic Allowed

```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te lsp (ingress)

```
show mpls rsvp-te lsp [ingress | <ingress_lsp_name> | ingress
<ingress_lsp_name> | ingress [destination | origin] <ipaddress>] {[all-
paths | detail] | summary | down-paths {detail}}
```

Description

Displays the ingress LSP information associated with RSVP-TE that is used to forward packets within the MPLS network.

Syntax Description

ingress	Displays only LSPs that originate from the switch into the MPLS network.
<ingress_lsp_name>	Displays only the named LSP that originates from the switch into the MPLS network.
destination <ipaddress>	Displays only those LSPs that terminate at the specified IP address.
origin <ipaddress>	Displays only those LSPs that originate at the specified IP address.
all-paths	Displays all paths for an LSP. The Active path has an asterisk (*) appended to the pathname.
detail	Displays the LSP information in verbose format.
summary	Displays only LSP summary statistics.
down-paths	Displays only paths with a “down” status.

Default

N/A.

Usage Guidelines

This command displays the LSP information associated with RSVP-TE that is used to forward packets within the MPLS network. If no options are specified, summary information for all RSVP-TE LSPs is displayed. The information included in the summary output will depend upon which, if any, additional keywords you use. If you specify only the `summary` keyword, the output statistics will include:

- LSP count.
- LSP direction relative to the MPLS network (i.e., ingress, egress, or transit).
- Number of LSPs enabled or disabled (ingress only).
- Number of LSPs with no configured path (ingress only).
- Primary/secondary path state (ingress only).

If the optional LSP name parameter is specified, the LSP information is displayed for the named LSP only. The LSPs displayed can be further qualified by the optional keyword `ingress`. This keyword qualifies the LSPs displayed from the perspective of the switch. Ingress LSPs identify LSPs that originate from the switch into the MPLS network.

The optional `destination` keyword limits the display to only those LSPs that terminate at the specified IP address. The optional `origin` keyword limits the display to only those LSPs that originate at the

specified IP address. The optional keyword `all-paths` displays all the paths for an LSP, not just the Active path. The Active path is shown with an asterisk (*) appended to the pathname. The `all-paths` keyword cannot be used with the `detail` option because the `detail` output shows all paths by default. The optional `down-paths` keyword displays only LSP paths whose status is “down.” Note that LSPs without a configured path will not be displayed with this keyword.

Specifying the optional `detail` keyword displays the LSP information in verbose format. Additional information displayed includes:

- Incoming and/or outgoing MPLS label.
- Configured profile.
- Cycle count.
- Receive packet and byte counters.
- Recorded routes.
- Number of VPLSs using ingress LSPs.
- Retry counts.
- Information about all configured primary and secondary paths.

Example

Use the following command to display information about a specific LSP:

```
* BD-10K.7 # show mpls rsvp-te lsp jefflsp1
```

Ingress LSP Name	Path Name	Destination	Transmit I/F	UpTm	Flags
jefflsp1	jeffpath1	11.100.100.204	n/a	0	--PR--IV

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

```
Summary of Ingress RSVP-TE LSPs named jefflsp1
  Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
  Ingress LSPs with no configured path  : 0
  Ingress LSP Paths (Up/Down)          : 2 (0/2)
```

Use the following command to display information about all RSVP-TE LSPs:

```
* BD-10K.5 # show mpls rsvp-te lsp
```

Ingress LSP Name	Path Name	Destination	Transmit I/F	UpTm	Flags
lsp598	path598	11.100.100.8	vlan1	3m	UEP--OIV

Egress LSP Name	Source IP	Destination	Receive I/F	UpTm
test_lsp5	11.100.100.8	11.100.100.20	vlan2	8m

Transit LSP Name	Source IP	Destination	Receive I/F	Transmit I/F	UpTm
lsp5a8	11.100.100.5	11.100.100.8	vlan1	vlan2	5m

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

Use the following command to display detailed information about a specific ingress LSP:

```
* BD-10K.7 # show mpls rsvp-te lsp ingress "lsp598" detail
```

```
Ingress LSP Name: lsp598
  Destination   : 11.100.100.8          Admin Status : Enabled
  IP Traffic    : Allow                 #VPLS Cfgd   : 0
  VPN Traffic   : Allow                 #VPLS In-Use : 0

  Path Name: path598
    Profile Name : prof598
    Tunnel ID    : 1                    Ext Tunnel ID : 11.100.100.20
    LSP ID       : 0                    State Changes  : 5
    Oper Status  : Enabled              Bandwidth Cfgd : False
    LSP Type     : Primary
    Activity     : Active
    Failures     : 2                    Retries-since last failure : 0
    Retries-Total : 12
    Rcv Label    : 0x0052e              UpTime        : 0d:0h:3m:44s
    Next Hop     : 11.121.96.5
    Tx I/F       : 11.121.96.20 - vlan1
    Record Route : Indx  IP Address
                  :    1  11.121.96.5
                  :    2  11.95.96.9
                  :    3  11.98.96.8
```

Use the following command to display detailed information about all paths for an LSP:

```
* BD-10K.5 # show mpls rsvp-te lsp jefflsp1 all-paths
```

Ingress LSP Name	Path Name	Destination	Transmit I/F	UpTm	Flags
jefflsp1	jeffpath0	11.100.100.204	n/a	0	--SR--IV
jefflsp1	jeffpath1	11.100.100.204	n/a	0	--PR--IV

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed
 (*) Indicates path is active

```
Summary of Ingress RSVP-TE LSPs named jefflsp1
  Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
  Ingress LSPs with no configured path : 0
  Ingress LSP Paths (Up/Down)          : 2 (0/2)
```

Use the following command to display detailed information about down paths for an LSP:

```
* BD-10K.5 # show mpls rsvp-te lsp ingress down-paths
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath0        11.100.100.204   n/a               0  --SR--IV
jefflsp1          jeffpath1        11.100.100.204   n/a               0  --PR--IV
jefflsp2          jeffpath2        11.100.100.203   n/a               0  -EPR-OIV
```

Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
 (R) Redundant Paths, (B) Bandwidth Requested, (O) ERO Specified,
 (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
 (v) VPN Assigned Traffic Allowed

```
Summary of Ingress RSVP-TE LSPs
  Ingress LSP Paths that are Down      : 3
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te neighbor

```
show mpls rsvp-te neighbor [{vlan} <vlan_name> | <ipaddress>] {detail}
```

Description

Displays all recognized RSVP-TE neighbors.

Syntax Description

vlan <vlan_name>	Displays only the neighbors for the specified VLAN.
<ipaddress>	Displays only the neighbor with the specified ipaddress.

Default

N/A.

Usage Guidelines

This command displays all recognized RSVP-TE neighbors. The IP address of each neighbor is displayed along with the VLAN name for the MPLS interface. For each neighbor, the following information is displayed:

- Number of RSVP-TE LSPs.
- Number of *hello* periods that have elapsed without receiving a valid *hello*.
- Remaining time before next *hello* is sent.
- Remaining time before next bundle message is sent.
- Neighbor up time.
- Neighbor supports RSVP hello.
- RSVP hello state.
- Neighbor supports refresh reduction.

If *vlan_name* is specified, only neighbors for the matching VLAN are shown. If *ipaddress* is specified, only the neighbor with that IP address is shown. If the *detail* keyword is specified, the information is shown in a verbose manner.

Example

The following command displays all recognized RSVP-TE neighbors:

```
* BD-10K.5 # show mpls rsvp-te neighbor
```

NeighborIP	VLAN Name	#LSPs	#Miss	NxtHello	NxtBundl	Flag	UpTm
11.121.96.5	vlan1	2	0	2870	0	UTH-	15m
11.122.96.8	vlan2	2	0	1770	0	UTH-	15m

Flags: (U) Hello Session Up, (T) Two Way Hello, (H) Neighbor Supports Hello,
(R) Neighbor Supports Refresh Reduction

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te path

```
show mpls rsvp-te path {<path_name>} {detail}
```

Description

Displays the configuration and usage information for MPLS RSVP-TE routed paths.

Syntax Description

<path_name>	Displays configuration and usage information for the specified MPLS RSVP-TE path.
detail	Displays the path information in verbose format.

Default

N/A.

Usage Guidelines

This command displays the configuration and usage information for MPLS RSVP-TE paths. Information is listed in tabular format and includes:

- Path name.
- Number of configured ERO objects.
- Number of LSPs configured to use this path.
- List of EROs and their type.

Specifying the optional `detail` keyword displays the path information in verbose format. If `detail` is specified, all LSPs that are configured to use the path are also displayed.

Example

The following command displays configuration and status information for the specified MPLS RSVP-TE paths:

```
* BD-10K.2 # show mpls rsvp-te path path598
Path Name          #LSP #ERO Ord#   ERO IP Netmask      Type
-----
path598            1    3   100 11.100.100.5/32   loose
                   200 11.100.100.9/32   loose
                   300 11.100.100.8/32   loose
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show mpls rsvp-te profile

```
show mpls rsvp-te profile {<profile_name>} {detail}
```

Description

Displays the configuration for the specified profile.

Syntax Description

<profile_name>	Displays configuration and usage information for the specified profile.
detail	Displays the profile information in verbose format.

Default

N/A.

Usage Guidelines

This command displays the configuration for the specified profile. All configured profile parameters are displayed in tabular format including: Peak burst parameters.

- Peak rate.
- Committed rate.
- Maximum burst size.
- Setup LSP priority.
- Hold LSP priority.
- Route recording status.
- Signaled MTU value.
- Number of LSPs using this profile.

If the `profile_name` argument is omitted, the profile parameter values for all profiles are displayed. Specifying the optional `detail` keyword displays the profile information in verbose format. When `detail` is specified, all LSPs that are configured to use this profile are also displayed.

Example

The following command displays configuration information for all defined profiles:

```
* BD-10K.13 # show mpls rsvp-te profile
Profile Name      Peak      Committed  Max Burst  SPri  HPri  RRO  MTU   #LSP
-----
default           0         0          0         7     0 Off i/f     0
prof598           0         0          0         7     0 On  1500     1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show vpls

```
show vpls {<vpls_name>} {detail}
```

Description

Displays VPLS configuration and status information.

Syntax Description

vpls_name	Displays information for the specified vpls.
detail	Displays additional information in comprehensive detail format.

Default

N/A.

Usage Guidelines

This command displays VPLS configuration and status information. The `show vpls` command (without specifying any optional parameters) displays all currently configured VPLS instances for the switch. Information displayed includes the:

- VPLS name
- IP addresses of peers
- Peer mode
- PW state
- Local egress VLAN or ports

If the name is specified, information is displayed for the specified VPLS. If the optional `detail` keyword is specified, additional information is displayed in comprehensive detail format. The detail format additionally displays transmit and receive packet and byte counts. The summarized list of VPLS instances is displayed in alphabetical order based on the `vpls_name`. Peers are displayed in the reverse of the order they were added.

Example

The following command shows detailed configuration and status information for the VPLS called "myvpls":

```
* BD-10000.1 # show vpls myvpls detail
      VPLS Name: myvpls                VPN ID: 64
      Admin State: Enabled              State: Enabled
      Source Address: 1.1.1.55          MTU: 1500
      Ethertype: 0x8100                .1q tag: exclude
      Services on: serv
```

```
Peer Name/IP: 1.1.1.202
Peer State: Up
Transport LSP: Unnamed LSP (best LSP to peer)
Next Hop IF: gtag103
  PW Uptime: 0d:4h:33m:58s      Transport LSP Tx Label: 0x80401
  Rx Label: 0x80402              Tx Label: 0x80402
  Rx Pkts: 0
  Rx Bytes: 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

show vpls peer

```
show vpls peer <ipaddress> {detail}
```

Description

Displays VPN VPLS information for the specified VPLS peer.

Syntax Description

ipaddress	Specifies the VPLS peer.

Default

N/A.

Usage Guidelines

This command displays L2 VPN VPLS information for the specified VPLS peer. The `ipaddress` parameter specifies the VPLS peer. Information displayed includes a list of all pseudo wires established to the peer, the pseudo wire status and PW ID, and information about each VPLS to which this peer belongs.

Example

The following command shows L2 VPN information for the specified VPLS peer:

```
show vpls peer 192.10.1.1
```

```
* BD-10000.2 # show vpls peer 1.1.1.202
VPLS Name      VPN ID  Flags  Services  Name      Peer IP      State  Flags
-----
myvpls         64      EAX-   serv      1.1.1.202  UP          C---
```

The following command shows detailed L2 VPN information for the specified VPLS peer:

```
VPLS Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
              (X) Exclude Tag, (T) Ethertype Configured
Peer Flags: (C) Core Peer, (S) Spoke Peer, (N) Named LSP Configured
```

```
* BD-10000.3 # show vpls peer 1.1.1.202 detail
VPLS Name: myvpls          VPN ID: 64
Source Address: 1.1.1.55    MTU: 1500
Services on: serv

Peer Name/IP: 1.1.1.202
Peer State: Up
```

```
Transport LSP: Unnamed LSP (best LSP to peer)
Next Hop IF: gtag103
PW Uptime: 0d:4h:37m:33s      Transport LSP Tx Label: 0x80401
Rx Label: 0x80402              Tx Label: 0x80402
Rx Pkts: 0
Rx Bytes: 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

traceroute mpls lsp

```
traceroute mpls lsp [<lsp_name> | any <host> | prefix <ipNetmask>]
{reply-mode [ip | mpls]} {{from <from>} {ttl <ttl>}}
{next-hop <hopaddress>}}
```

Description

Traces the path an LSP takes for the specified FEC.

Syntax Description

lsp_name	Specifies the LSP on which to send the MPLS echo request.
any	Allows the echo request to be sent over any available LSP.
<host>	Specifies the FEC using an <i>ipaddress</i> or <i>hostname</i> .
prefix	Specifies a prefix.
<ipNetmask>	Specifies the prefix address.
reply-mode	Specifies the return path for the MPLS echo response.
ip	Specifies the return path for the MPLS echo response. Routes back to the sender via an IPv4 packet with the IP Router Alert option. This is the default mode.
mpls	Specifies the return path for the MPLS echo response. Switches the echo reply back to the sender via an IPv4 packet over an LSP.
from <from>	Specifies the IP address to be used as the source address in the MPLS echo request.
ttl	Specifies the maximum TTL value in the MPLS echo request packet.
<ttl>	Specifies the maximum TTL value in seconds. The range is from 1 - 30. The default is 30.
next-hop	Routes the MPLS echo request along the LSP that traverses the specified node.
<hopaddress>	Specifies the next-hop address.

Default

The maximum time-to-live value is 30 seconds.

Usage Guidelines

This command traces the path an LSP takes for the specified FEC. The `traceroute` command, with the `mpls` keyword option, works by repeatedly sending an MPLS echo request (or "MPLS Ping"). The TTL value is incremented for each successive MPLS echo request sent. The sending LSR waits 5 seconds before sending the next MPLS echo request. This operation continues until either the egress LSR for the FEC is reached, the maximum TTL value is reached, or the operation is interrupted. For each response received, the following information is displayed on the console:

- IP address of the replying LSR
- Return code
- Indication of an MPLS echo reply timeout if no response was received

The FEC can be specified using the *ipaddress* or *hostname* via the `host` parameter. If the optional `next-hop` is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This

option is useful for tracing a specific LSP when multiple LSPs exist to the specified FEC. The `lsp` keyword may be used to specify a named LSP to trace. The selected LSP is specified by the `lsp_name` parameter. The `any` keyword indicates that the switch can trace any available LSP to the specified host.

The optional `reply-mode` keyword is used to specify the return path for the MPLS echo response. When the `ip` option is specified, the MPLS echo reply is routed back to the sender via an IPv4 packet with the IP Router Alert option. When the `mpls` option is specified, the MPLS echo reply is switched back to the sender via an IPv4 packet over an LSP. If the `reply-mode` is not specified, the `ip` reply mode is used.

The optional `ttl` keyword optionally specifies the maximum TTL value in the MPLS echo request packet. The TTL value is decremented along the path at each LSR. When the TTL value reaches zero, the packet is not forwarded but is redirected to the control plane for additional processing. When the TTL value specified by the `ttl` parameter is exceeded, the MPLS LSP traceroute command ends. The valid `ttl` range is 1 to 30 seconds with a default value of 30 seconds.

The `from` keyword is used to specify the source IP address used in the MPLS echo request. This is the IP address used by the target LSR will use to send the MPLS echo reply. If not specified, the OSPF router ID is used.

Example

The following example shows a sample display for the `traceroute` command:

```
BD-10K.5 # traceroute mpls lsp prefix 11.100.100.10/32
traceroute to 11.100.100.10, 30 hops max
 1  11.100.100.8                    5 ms          5 ms          2 ms
 2  11.100.100.10                  2 ms          1 ms          2 ms
BD-10K.6 #
BD-10K.6 # traceroute mpls lsp lsp598
traceroute to lsp598, 30 hops max
 1  11.100.100.5                   6 ms          1 ms          5 ms
 2  11.100.100.9                   3 ms          2 ms          2 ms
 3  11.100.100.8                   3 ms          4 ms          3 ms
BD-10K.7 #
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

unconfigure mpls

```
unconfigure mpls
```

Description

Resets MPLS configuration parameters to the default settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command deletes all VLANs from MPLS and resets all MPLS configuration parameters to their default values. The parameters that are reset include the LSR ID, all LDP-specific settings, all RSVP-TE-specific settings, all RSVP-TE reserved bandwidth, and all EXP Qos Profile mappings. MPLS must be disabled to unconfigure MPLS.



NOTE

MPLS must be disabled to globally unconfigure MPLS.

Example

The following command resets MPLS configuration parameters to the default settings:

```
unconfigure mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

unconfigure mpls exp examination

```
unconfigure mpls exp examination [all | {{value} <value>}]
```

Description

Resets the QoS profile assigned to the EXP value back to the default QoS profile.

Syntax Description

Deny

value <value>	Specifies the EXP value whose QoS profile will be reset.

Default

The QoS profile matches the EXP value + 1

Usage Guidelines

This command resets the QoS profile assigned to the EXP value (defined by the `value` keyword and argument) back to the default QoS profile. If the `all` option is specified, all EXP values are reset back to their default QoS profiles. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance.

Example

Use the following command to restore the QoS profile of EXP value 5 to its default setting:

```
unconfigure mpls exp examination value 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

unconfigure mpls exp replacement

```
unconfigure mpls exp replacement [all | {{qosprofile} <qosprofile>}]
```

Description

Resets the EXP value assigned to the QoS profile back to the default EXP value.

Syntax Description

deny	
qosprofile	Specifies the QoS profile whose EXP replacement value will be unconfigured.
<qosprofile>	

Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command resets the EXP value assigned to the QoS defined by `qosprofile` back to the default EXP value. If the `all` option is specified, all QoS profiles are reset back to their default EXP values. By default, the EXP value matches the QoS profile - 1. That is, QoS profile `qp1` is mapped to EXP value of 0, QoS profile `qp2` is mapped to EXP value of 1, etc. This configuration has switch-wide significance.

Example

Use the following command to restore all EXP values to their default setting:

```
unconfigure mpls exp replacement all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

unconfigure mpls vlan

```
unconfigure mpls {{vlan} <vlan_name>}
```

Description

Resets MPLS configuration parameters to the default settings. This command does not delete the VLAN from MPLS.

Syntax Description

vlan <vlan_name>	Specifies the VLANs for which MPLS will be unconfigured.

Default

N/A.

Usage Guidelines

This command resets all MPLS configuration parameters for the specified VLAN to their default values. It does not delete the VLAN from MPLS. These parameters include the enable state for LDP and RSVP-TE, the bandwidth reserved for RSVP-TE LSPs, RSVP-TE timers, and the RSVP-TE metric. MPLS does not have to be disabled to unconfigure a specific VLAN.

Example

The following command resets MPLS configuration parameters to the default settings for a single VLAN:

```
unconfigure mpls vlan boone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

unconfigure vpls dot1q ethertype

```
unconfigure vpls <vpls_name> dot1q ethertype
```

Description

Unconfigures the ethertype setting for the VPLS specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)

Default

N/A

Usage Guidelines

This command unconfigures the ethertype setting for the VPLS specified by `vpls_name`. The setting is changed back to the value displayed in the `show dot1q` command.

Example

The following command changes the ethertype setting for the specified VPLS to the value displayed in the `show dot1q` command:

```
unconfigure vpls my_vpls dot1q ethertype
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the BlackDiamond 10808 with an MSM-1XL and BlackDiamond 12000 series switches with an enabled MPLS feature pack.

Extreme Networks switches provide full Layer 3, IP unicast routing. They exchange routing information with other routers on the network using either the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. The switches dynamically build and maintain routing tables and determine the best path for each of its routes.

Each host that uses the IP unicast routing functionality of the switch must have a unique IP address assigned. In addition, the default gateway assigned to the host must be the IP address of the router interface.

The routing software and hardware directs IP traffic between router interfaces. A router interface is simply a VLAN that has an IP address assigned to it.

As you create VLANs with IP addresses belonging to different IP subnets, you can also choose to route between the VLANs. The VLAN switching and IP routing functions occur within the switch.

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the same IP subnet on different VLANs.

The Extreme Networks switch maintains an IP routing table for network routes and host routes. The table is populated from the following sources:

- Dynamically, by way of routing protocol packets or by ICMP redirects exchanged with other routers
- Statically, by way of routes entered by the administrator
 - Default routes, configured by the administrator
 - Locally, by way of interface addresses assigned to the system
 - By other static routes, as configured by the administrator

Dynamic routes are typically learned by way of RIP or OSPF. Routers that use RIP or OSPF exchange information in their routing tables in the form of advertisements. Using dynamic routes, the routing table contains only networks that are reachable.

Dynamic routes are aged out of the table when an update for the network is not received for a period of time, as determined by the routing protocol.

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers.

Static routes can also be used for security reasons, to control which routes you want advertised by the router. Static routes are never aged out of the routing table.

A static route must be associated with a valid IP subnet. An IP subnet is associated with a single VLAN by its IP address and subnet mask. If the VLAN is subsequently deleted, the static route entries using that subnet must be deleted manually.

When there are multiple, conflicting choices of a route to a particular destination, the router picks the route with the longest matching network mask. If these are still equal, the router picks the route using the following criteria (in the order specified):

- Directly attached network interfaces
- ICMP redirects
- Static routes
- Directly attached network interfaces that are not active

If you define multiple default routes, the route that has the lowest metric is used. If there are multiple default routes that have the same lowest metric, the system picks one of the routes with the lowest gateway IP addresses.

You can also configure *blackhole* routes—traffic to these destinations is silently dropped.

Internet Control Message Protocol (ICMP) is used to transmit information needed to control IP traffic. It is used mainly to provide information about routes to destination addresses. ICMP redirect messages inform hosts about more accurate routes to other systems, whereas ICMP unreachable messages indicate problems with a route.

Additionally, ICMP can cause TCP connection to terminate gracefully if the route becomes unavailable.

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets. This feature can be used in various applications, including DHCP services between Windows NT servers and clients running Windows 95.

Proxy Address Resolution Protocol (ARP) was first developed so that ARP-capable devices could respond to ARP Request packets on behalf of ARP-incapable devices. Proxy ARP can also be used to achieve router redundancy and simplify IP client configuration. The Extreme Networks switch supports proxy ARP for this type of network configuration.

After IP ARP has been configured, the system responds to ARP Requests on behalf of the device, as long as the following conditions are satisfied:

- The valid IP ARP Request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The proxy ARP table entry indicates that the system should always answer this ARP Request, regardless of the ingress VLAN (the `always` parameter must be applied).

After all the proxy ARP conditions have been met, the switch formulates an ARP Response using the configured MAC address in the packet.

In some networks, it is desirable to configure the IP host with a wider subnet than the actual subnet mask of the segment. Proxy ARP can be used so that the router answers ARP Requests for devices outside of the subnet. As a result, the host communicates as if all devices are local. In reality, communication with devices outside of the subnet are proxied by the router.



NOTE

Beginning with ExtremeXOS version 12.0, support for route handling optimization/compression was added to allow you to reduce the number of routes that are being installed in hardware by installing less specific ones in the case of overlapping routes.

clear iparp

```
clear iparp {<ip_address> {vr <vr_name>} | vlan <vlan_name>}
```

Description

Removes dynamic entries in the IP ARP table.

Syntax Description

ip_address	Specifies an IP address.
vlan_name	Specifies a VLAN name.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Permanent IP ARP entries are not affected.

This command is specific to a single virtual router, and it applies to the current virtual router if you do not specify a virtual router.

Example

The following command removes a dynamically created entry from the IPARP table:

```
clear iparp 10.1.1.5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootprelay add

```
configure bootprelay add <ip_address> {vr <vrid>}
```

Description

Configures the addresses to which BOOTP requests should be directed.

Syntax Description

ip_address	Specifies an IP address.
vrid	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

- 1 Configure VLANs and IP unicast routing.
- 2 Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootprelay add <ip_address>
```

- 3 Enable the DHCP or BOOTP relay function, using the following command:

```
enable bootprelay
```

Example

The following command configures BOOTP requests to be directed to 123.45.67.8:

```
configure bootprelay add 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootprelay delete

```
configure bootprelay delete [<ip_address> | all] {vr <vrid>}
```

Description

Removes one or all IP destination addresses for forwarding BOOTP packets.

Syntax Description

ip_address	Specifies an IP address.
vrid	Specifies a VR name.
all	Specifies all IP address entries.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command removes the destination address:

```
configure bootprelay delete 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootprelay dhcp-agent information check

```
configure bootprelay dhcp-agent information check
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To disable this check, use the following command:

```
unconfigure bootprelay dhcp-agent information check
```

Example

The following command configures the DHCP relay agent option check:

```
configure bootprelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure bootrelay dhcp-agent information option

```
configure bootrelay dhcp-agent information option
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward DHCP or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

- 1 Configure VLANs and IP unicast routing.
- 2 Enable the DHCP or BOOTP relay function, using the following command:

```
enable bootrelay {vr <vrid>}
```
- 3 Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootrelay add <ip_address> {vr <vrid>}
```
- 4 Configure the DHCP relay agent option (option 82), using the following command:

```
configure bootrelay dhcp-agent information option
```

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure bootrelay dhcp-agent information option
```

Example

The following command configures the DHCP relay agent option:

```
configure bootrelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure bootprelay dhcp-agent information policy

```
configure bootprelay dhcp-agent information policy [drop | keep | replace]
```

Description

Configures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

Replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure bootprelay dhcp-agent information policy keep
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure iparp add

```
configure iparp add <ip_addr> {vr <vr_name>} <mac>
```

Description

Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.

Syntax Description

ip_addr	Specifies an IP address.
mac	Specifies a MAC address.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Add a permanent IP ARP entry to the system. The `ip_address` is used to match the IP interface address to locate a suitable interface.

Example

The following command adds a permanent IP ARP entry to the switch for IP address *10.1.2.5*:

```
configure iparp add 10.1.2.5 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp add proxy

```
configure iparp add proxy [<ipNetmask> | <ip_addr> {<mask>}] {vr <vr_name>}
{<mac>} {always}
```

Description

Configures the switch to respond to ARP Requests on behalf of devices that are incapable of doing so.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
mac_address	Specifies a MAC address.
always	Specifies that the switch responds regardless of the VLAN that the request arrives from.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

When `mask` is not specified, an address with the mask 255.255.255.255 is assumed. When `mac_address` is not specified, the MAC address of the switch is used in the ARP Response. When `always` is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.

After IP ARP is configured, the system responds to ARP Requests on behalf of the device as long as the following conditions are satisfied:

- The valid IP ARP Request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The source IP address is not on the same subnet as the target address (unless the `always` flag is set).

After all the proxy ARP conditions have been met, the switch formulates an ARP Response using the configured MAC address in the packet.

The default maximum number of proxy entries is 256, but can be increased to 4096 by using the following command:

```
configure iparp max_proxy_entries <max_proxy_entries>
```

Example

The following command configures the switch to answer ARP Requests for all devices with the address range of 100.101.45.1 to 100.101.45.255:

```
configure iparp add proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp delete

```
configure iparp delete <ip_addr> {vr <vr_name>}
```

Description

Deletes an entry from the ARP table. Specify the IP address of the entry.

Syntax Description

ip_addr	Specifies an IP address.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Removes any IP ARP entry (dynamic or permanent) from the table. The `ip_address` is used to match the IP interface address to locate a suitable interface.

Example

The following command deletes an IP address entry from the ARP table:

```
configure iparp delete 10.1.2.5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp delete proxy

```
configure iparp delete proxy [[<ipNetmask> | <ip_addr> {<mask>}]] {vr  
<vr_name>} | all]
```

Description

Deletes one or all proxy ARP entries.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
all	Specifies all ARP entries.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

When the `mask` is not specified, the software will assume a host address (that is, a 32-bit mask).

Example

The following command deletes the IP ARP proxy entry *100.101.45.0/24*:

```
configure iparp delete proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp max_entries

```
configure iparp max_entries <max_entries>
```

Description

Configures the maximum allowed IP ARP entries.

Syntax Description

max_entries	Specifies a number of maximum IP ARP entries.

Default

4096.

Usage Guidelines

Range: 1 - 20480. The maximum IP ARP entries include dynamic, static, and incomplete IP ARP entries.

Example

The following command sets the maximum IP ARP entries to 2000 entries:

```
configure iparp max-entries 2000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp max_pending_entries

```
configure iparp max_pending_entries <max_pending_entries>
```

Description

Configures the maximum allowed incomplete IP ARP entries.

Syntax Description

max_pending_entries	Specifies a number of maximum IP ARP entries.

Default

256.

Usage Guidelines

Range: 1 - 4096.

Example

The following command sets the maximum pending IP ARP entries to 500 entries:

```
configure iparp max_pending_entries 500
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp max_proxy_entries

```
configure iparp max_proxy_entries <max_proxy_entries>
```

Description

Configures the maximum allowed IP ARP proxy entries.

Syntax Description

max_proxy_entries	Specifies maximum number of IP ARP proxy entries.

Default

256.

Usage Guidelines

Range: 0 - 4096.

Example

The following command sets the maximum IP ARP proxy entries to 500 entries:

```
configure iparp max_proxy_entries 500
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure iparp timeout

```
configure iparp timeout {vr <vr_name>} <minutes>
```

Description

Configures the IP ARP timeout period.

Syntax Description

vr_name	Specifies which virtual router IP ARP setting to change.
minutes	Specifies a time in minutes.

Default

20 minutes.

Usage Guidelines

The range is 0-32,767. A setting of 0 disables timeout.

When the switch learns an ARP entry, it begins the timeout for that entry. When the timer reaches 0, the entry will be aged out, unless IP ARP refresh is enabled. If ARP refresh is enabled, the switch will ARP for the address before the timer has expired. If the switch receives a response, it will reset the timer for that address.

Example

The following command sets the IP ARP timeout period to 10 minutes:

```
configure iparp timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ipforwarding originated-packets

```
configure ipforwarding originated-packets [ require-ipforwarding | dont-require-ipforwarding ]
```

Description

Configures the ipforwarding behavior.

Syntax Description

This command has no arguments or variables.

Default

dont-require-ipforwarding.

Usage Guidelines

You can configure IP forwarding to be either required or not required. If not configured, the default is for forwarding to be not required.

Example

The following command configures the ipforwarding:

```
configure ipforwarding originated-packets require-ipforwarding
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure iproute add

```
configure iproute add [<ipNetmask> | <ip_addr> <mask>] <gateway> <metric>
{vr <vrname>} {multicast-only | unicast-only}
```

Description

Adds a static route to the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
metric	Specifies a cost metric.
vrname	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 for mask to indicate a host entry.

Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute add blackhole

```
configure iproute add blackhole [<ipNetmask> | <ipaddress> <mask>] {vr
<vrname>} {multicast-only | unicast-only}
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
vrname	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry configures packets with the specified destination IP subnet to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination IP subnet must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Example

The following command adds a blackhole address to the routing table for packets with a destination address of 100.101.145.4:

```
configure iproute add blackhole 100.101.145.4/32
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute add blackhole default

```
configure iproute add blackhole {ipv4} default {vr <vrname>} {multicast-  
only | unicast-only}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

vr_name	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is 0.0.0.0.

Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add default

```
configure iproute add default <gateway> {vr <vrname>} {<metric>}  
{multicast-only | unicast-only}
```

Description

Adds a default gateway to the routing table.

Syntax Description

gateway	Specifies a VLAN gateway
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. Use the `unicast-only` or `multicast-only` options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 123.45.67.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete

```
configure iproute delete [<ipNetmask> | <ipaddress> <mask>] <gateway> {vr  
<vrname>}
```

Description

Deletes a static address from the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following command deletes an address from the gateway:

```
configure iproute delete 10.101.0.0/24 10.101.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole

```
configure iproute delete blackhole [<ipNetmask> | <ipaddress> <mask>] {vr  
<vrname>}
```

Description

Deletes a blackhole address from the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a netmask.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command removes a blackhole address from the routing table:

```
configure iproute delete blackhole 100.101.145.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole default

```
configure iproute delete blackhole default {vr <vrname>}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

vrname	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete default

```
configure iproute delete default <gateway> {vr <vrname>}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

gateway	Specifies a VLAN gateway.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default gateway:

```
configure iproute delete default 123.45.67.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute priority

```
configure iproute {ipv4} priority [rip | blackhole | bootp | ebgp | ibgp |
icmp | static | ospf-intra | ospf-inter | ospf-as-external | ospf-extern1 |
ospf-extern2] <priority>
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

rip	Specifies RIP.
bootp	Specifies BOOTP.
ebgp	Specifies E-BGP routes
ibgp	Specifies I-BGP routes
icmp	Specifies ICMP.
blackhole	Specifies the blackhole route.
static	Specifies static routes.
ospf-intra	Specifies OSPFIntra routing.
ospf-inter	Specifies OSPFInter routing.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
priority	Specifies a priority number in the range of 11 to 65534.

Default

Table 26 lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 26: Relative route priorities

Route Origin	Priority
Direct	10
Blackhole	50
Static	1100
ICMP	1200
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPFAsExt	3100
OSPF External 1	3200
OSPF External 2	3300
BOOTP	5000

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.

**NOTE**

The priority for a blackhole route can not overlap with the priority of any other route origin.

Example

The following command sets IP route priority for static routing to 1200:

```
configure iproute priority static 1200
```

History

This command was first available in ExtremeXOS 10.1.

The route priority restrictions were added in ExtremeXOS 11.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute sharing max-gateways

```
configure iproute sharing max-gateways <max_gateways>
```

Description

Configures the maximum number of equal-cost multipath (ECMP) gateways per destination supported in hardware.

Syntax Description

max-gateways	Specifies the maximum number of ECMP gateways per destination supported in the switch hardware. The only values allowed are 4 and 8.

Default

- BlackDiamond 8800 family of switches, SummitStack, and Summit family switch only—4
- BlackDiamond 10808 and BlackDiamond 12800 series switches only—8 (not configurable)

Usage Guidelines

You must save the configuration and reboot the switch for the new value to take effect. To see the current and configured value, use the following command:

```
show ipconfig
```

Example

The following command changes the maximum number of ECMP gateways per destination supported in hardware to eight:

```
configure iproute sharing max-gateways 8
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available only on the SummitStack, Summit family switches, and the BlackDiamond 8800 family of switches.

configure irdp

```
configure irdp [multicast | broadcast | <mininterval> <maxinterval>  
<lifetime> <preference>]
```

Description

Configures the destination address of the router advertisement messages.

Syntax Description

multicast	Specifies multicast setting.
broadcast	Specifies broadcast setting.
mininterval	Specifies the minimum time between advertisements.
maxinterval	Specifies the maximum time between advertisements. Default is 600.
lifetime	Specifies the lifetime of the advertisement. Default is 1800.
preference	Specifies the router preference level. Default is 0.

Default

Broadcast (255.255.255.255). The default mininterval is 450.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch will send out IP packets at the intervals specified identifying itself as a default router. IRDP enabled client machines will use this information to determine which gateway address to use for routing data packets to other networks.

Example

The following command sets the address of the router advertiser messages to multicast:

```
configure irdp multicast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan add secondary-ipaddress

```
configure vlan <vlan_name> add secondary-ipaddress [<ipaddress> {<netmask>}  
| <ipNetmask>]
```

Description

Configures secondary IP addresses on a VLAN to support multinetting.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IP address.
netmask	Specifies a network mask.
ipNetmask	Specifies an IP address with network mask.

Default

N/A.

Usage Guidelines

Adding a secondary IP address to a VLAN enables multinetting. Secondary addresses are added to support legacy stub IP networks.

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Delete secondary address with the following command:

```
configure vlan <vlan_name> delete secondary-ipaddress [<ipaddress> | all]
```

Example

The following command configures the VLAN *multi* to support the 10.1.1.0/24 subnet in addition to its primary subnet:

```
configure vlan multi add secondary-ipaddress 10.1.1.1/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan delete secondary-ipaddress

```
configure vlan <vlan_name> delete secondary-ipaddress [<ipaddress> | all]
```

Description

Removes secondary IP addresses on a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IP address.
all	Specifies all secondary IP addresses.

Default

N/A.

Usage Guidelines

Delete secondary IP addresses that were added to support multinetting.

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Use the `all` keyword to delete all the secondary IP addresses from a VLAN.

Example

The following command removes the 10.1.1.0 secondary IP address from the VLAN *multi*:

```
configure vlan multi delete secondary-ipaddress 10.1.1.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan udp-profile

```
configure vlan <vlan_name> udp-profile [<profilename> | none]
```

Description

Associates a UDP Forwarding profile to a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
profilename	Specifies a policy file to use for the UDP forwarding profile.
none	Removes any UDP forwarding profile from the VLAN.

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

You can apply a UDP forwarding policy only to an L3 VLAN (a Vlan having at least one IP address configured on it). If there is no IP address configured on the VLAN, then the command will be rejected.

A UDP forwarding policy must contain only the following attributes. Unrecognized attributes will be ignored.

- Match attributes
 - Destination UDP Port Number (destination-port)
 - Source IP address (source-ipaddress)
- Action modified (set) attributes
 - Destination IP address (destination-ipaddress)
 - VLAN name (vlan)

Policy files used for UDP forwarding are processed differently from standard policy files. Instead of terminating when an entry's match clause becomes true, each entry in the policy file will be processed and the corresponding action will be taken for each true match clause.

For example, if the following policy file is used as a UDP forwarding profile, any packets destined for UDP port 67 will be sent to IP address 20.0.0.5 AND flooded to VLAN *to7*:

```
entry one {
  if match all {
    destination-port 67 ;
  } then {
    destination-ipaddress 20.0.0.5 ;
  }
}
```

```
entry two {
  if match all {
    destination-port 67 ;
  } then {
    vlan "to7" ;
  }
}
```

If you include more than one VLAN set attribute or more than one destination-ipaddress set attribute in one policy entry, the last one will be accepted and the rest ignored. Therefore, you can have two valid set statements in each entry of a UDP forwarding policy; one a destination-ipaddress and one a VLAN. ExtremeXOS currently allows a maximum of eight entries in a UDP forwarding policy, so you can define a maximum of sixteen destinations for one inbound broadcast UDP packet: eight IP address and eight VLANs.

**NOTE**

It is strongly advised not to have more than eight entries in a UDP forwarding profile. The UDP forwarding module will process those entries even if the entries do not contain any attributes for UDP forwarding. Having more than eight entries will drastically reduce the performance of the system. If the inbound UDP traffic rate is very high, having more than eight entries could cause the system to freeze or become locked.

**NOTE**

If you rename a VLAN referred to in your UDP forwarding profile, you must manually edit the policy to reflect the new name, and refresh the policy.

You can also validate whether the UDP profile has been successfully associated with the VLAN by using the command `show policy {<policy-name> | detail}`. UDP forwarding is implemented as part of the netTools process, so the command will display netTools as a user of the policy.

Example

The following command associates the UDP forwarding profile `port123_to_corporate` to the VLAN `to-sales`:

```
configure vlan to-sales udp-profile port123_to_corporate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable bootp vlan

```
disable bootp vlan [<vlan> | all]
```

Description

Disables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

vlan	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
disable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable bootprelay

```
disable bootprelay {vr <vrid>}
```

Description

Disables the BOOTP relay function.

Syntax Description

vrid	Specifies the virtual router to be disabled.

Default

Disabled.

Usage Guidelines

This command can disable the BOOTP relay functionality for a particular virtual router, or all of them. If you use the command without specifying a virtual router, the functionality is disabled for all virtual routers.

Example

The following command disables the forwarding of BOOTP requests:

```
disable bootprelay
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp address-mask

```
disable icmp address-mask {vlan <name>}
```

Description

Disables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Disables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
disable icmp address-mask vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The default was changed to disabled in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable icmp parameter-problem

```
disable icmp parameter-problem {vlan <name>}
```

Description

Disables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
disable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp port-unreachables

```
disable icmp port-unreachables {vlan <name>}
```

Description

Disables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables ICMP port unreachable messages on VLAN *accounting*:

```
disable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp redirects

```
disable icmp redirects {vlan <name>}
```

Description

Disables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP redirects (Type 5) to hosts who direct routed traffic to the switch where the switch detects that there is another router in the same subnet with a better route to the destination.

Example

The following command disables ICMP redirects from VLAN *accounting*:

```
disable icmp redirects vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp time-exceeded

```
disable icmp time-exceeded {vlan <name>}
```

Description

Disables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
disable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp timestamp

```
disable icmp timestamp {vlan <name>}
```

Description

Disables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP timestamp response on VLAN *accounting*:

```
disable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp unreachable

```
disable icmp unreachable {vlan <name>}
```

Description

Disables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP unreachable messages on all VLANs:

```
disable icmp unreachable
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp useredirects

```
disable icmp useredirects
```

Description

Disables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

If the switch has a route to a destination network, the switch will use that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router will send an ICMP redirect message to the first router. If ICMP useredirects is disabled, the switch will disregard these messages, and continue to send the packets to the second router.

Example

The following command disables the changing of routing table information:

```
disable icmp useredirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable iparp checking

```
disable iparp {vr <vr_name>} checking
```

Description

Disable checking if the ARP Request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

vr_name	Specifies a virtual router.

Default

Enabled.

Usage Guidelines

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command disables IP ARP checking:

```
disable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable iparp refresh

```
disable iparp {vr <vr_name>} refresh
```

Description

Disables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

vr_name	Specifies a virtual router.

Default

Enabled.

Usage Guidelines

The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count layer 2 switching only environment.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command disables IP ARP refresh:

```
disable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable ipforwarding

```
disable ipforwarding {broadcast | ignore-broadcast | fast-direct-broadcast}
{vlan <vlan_name>}
```

Description

Disables routing (or routing of broadcasts) for one or all VLANs. If no argument is provided, disables routing for all VLANs.

Syntax Description

broadcast	Specifies broadcast IP forwarding.
ignore-broadcast	Specifies to forward IP broadcast packets in the hardware but not send a copy to the CPU. (Available on the BlackDiamond 10808 switch only.)
fast-direct-broadcast	Specifies to forward IP broadcast packets in the hardware and send a copy to the CPU. (Available on the BlackDiamond 10808 switch only.)
vlan_name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Disabling IP forwarding also disables broadcast forwarding. Broadcast forwarding can be disabled without disabling IP forwarding. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default.

Other IP related configuration is not affected.

Example

The following command disables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
disable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The ignore-broadcast and fast-direct-broadcast keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable iproute compression

```
disable iproute {ipv4} compression {vr <vrname>}
```

Description

Disables IP route compression.

Syntax Description

ipv4	Disable IP route compression algorithm for IPv4 address family. Default is IPv4.
vrname	Virtual router name for which the IP route compression is being disabled. If the virtual router name is not specified, the route compression will be disabled for the VR context from which CLI command is being issued.

Default

Disabled.

Usage Guidelines

Disables IP route compression for a specified address family and virtual router.

Example

```
disable iproute compression
```

By default, disables route compression for the IPv4 family for the VR context from which the CLI command is being used.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable ip-option loose-source-route

```
disable ip-option loose-source-route
```

Description

Disables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets with the IP option for loose source routing turned on. Packets with the loose source route option enabled will be dropped by the switch.

Example

The following command disables processing of the loose source route IP option:

```
disable ip-option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option record-route

```
disable ip-option record-route
```

Description

Disables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from adding itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command disables processing of the record route IP option:

```
disable ip-option record-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option record-timestamp

```
disable ip-option record-timestamp
```

Description

Disables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from adding a timestamp into the IP options header when it receives a packet with the record timestamp IP option.

Example

The following command disables processing of the record timestamp IP option:

```
disable ip-option record-timestamp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option strict-source-route

```
disable ip-option strict-source-route
```

Description

Disables processing the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets that have the strict source routing IP option turned on. The switch will drop packets that have the strict source routing IP option enabled.

Example

The following command disables processing of the strict source route IP option:

```
disable ip-option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option router-alert

```
disable ip-option router-alert
```

Description

Disables processing of the router alert IP option in IPv4 packet headers.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables processing of the router alert IP option:

```
disable ip-option router-alert
```

History

This command was first available in EXOS 10.1.

Platform Availability

This command is available on all platforms.

disable irdp

```
disable irdp {vlan <name>}
```

Description

Disables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command disables IRDP on VLAN *accounting*:

```
disable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable udp-echo-server

```
disable udp-echo-server {vr <vrid>}
```

Description

Disables UDP echo server support.

Syntax Description

vr	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving end.

Example

The following command disables UDP echo server support:

```
disable udp-echo-server
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable bootp vlan

```
enable bootp vlan [<vlan> | all]
```

Description

Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

vlan	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
enable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable bootprelay

```
enable bootprelay {vr <vrid>}
```

Description

Enables the BOOTP relay function.

Syntax Description

vrid	Specifies the virtual router to be disabled.

Default

Disabled. The default virtual router is *VR-Default*.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets. This feature can be used in various applications, including DHCP services between Windows NT servers and clients running Windows 95.

To configure the relay function, follow these steps:

- 1 Configure VLANs and IP unicast routing.
- 2 Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootprelay add <ip_address>
```

- 3 Enable the DHCP or BOOTP relay function, using the following command:

```
enable bootprelay
```

Example

The following command enables the forwarding of BOOTP requests:

```
enable bootprelay
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp address-mask

```
enable icmp address-mask {vlan <name>}
```

Description

Enables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Enables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
enable icmp address-mask vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The default was changed to disabled in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable icmp parameter-problem

```
enable icmp parameter-problem {vlan <name>}
```

Description

Enables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
enable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp port-unreachables

```
enable icmp port-unreachables {vlan <name>}
```

Description

Enables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables ICMP port unreachable messages on VLAN *accounting*:

```
enable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp redirects

```
enable icmp redirects {vlan <name>}
```

Description

Enables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is in routing mode.

ICMP redirects are used in the situation where there are multiple routers in the same subnets. If a host sends a packet to its default gateway, the gateway router will look at its route table to find the best route to the destination. If it sees that the best route is through a router in the same subnet as the originating host, the switch will send an ICMP redirect (type 5) message to the host that originated the packet, telling it to use the other router with the better route. The switch will also forward the packet to the destination.

Example

The following command enables the generation of ICMP redirect messages on all VLANs:

```
enable icmp redirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp time-exceeded

```
enable icmp time-exceeded {vlan <name>}
```

Description

Enables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
enable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp timestamp

```
enable icmp timestamp {vlan <name>}
```

Description

Enables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP timestamp response on VLAN *accounting*:

```
enable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp unreachable

```
enable icmp unreachable {vlan <name>}
```

Description

Enables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP unreachable messages on all VLANs:

```
enable icmp unreachable
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp useredirects

```
enable icmp useredirects
```

Description

Enables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the switch has a route to a destination network, the switch will use that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router will send an ICMP redirect message to the originating router. If ICMP useredirects is enabled, the switch will add a route to the destination network using the third router as the next hop and start sending the packets to the third router.

Example

The following command enables the modification of route table information:

```
enable icmp useredirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable iparp checking

```
enable iparp {vr <vr_name>} checking
```

Description

Enables checking if the ARP Request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

vr_name	Specifies a virtual router.

Default

Enabled.

Usage Guidelines

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command enables IP ARP checking:

```
enable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable iparp refresh

```
enable iparp {vr <vr_name>} refresh
```

Description

Enables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

vr_name	Specifies a virtual router.

Default

Enabled.

Usage Guidelines

IP ARP refresh is used so that the switch does not have to wait for the ARP timer for a host to expire before re-ARPing for the host. If ARP refresh is enabled, the switch will re-ARP for the host at 3/4th of the configured ARP timer.

For example: If the ARP timeout is set to 20 minutes, the switch will attempt to re-ARP for the host when the host entry is at 15 minutes. If the host replies, the ARP entry is reset back to 0, and the timer starts again.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command enables IP ARP refresh:

```
enable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable ipforwarding

```
enable ipforwarding {ipv4 | broadcast | ignore-broadcast | fast-direct-  
broadcast} {vlan <vlan_name>}
```

Description

Enables IPv4 routing or IPv4 broadcast forwarding for one or all VLANs. If no argument is provided, enables IPv4 routing for all VLANs that have been configured with an IP address on the current virtual router.

Syntax Description

ipv4	Specifies IPv4 forwarding
broadcast	Specifies broadcast IP forwarding.
ignore-broadcast	Specifies to forward IP broadcast packets in the hardware but not send a copy to the CPU. (Available on the BlackDiamond 10808 switch only.)
fast-direct-broadcast	Specifies to forward IP broadcast packets in the hardware and send a copy to the CPU. (Available on the BlackDiamond 10808 switch only.)
vlan_name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

IP forwarding must first be enabled before IP broadcast forwarding can be enabled. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default. Currently, Extreme Networks switches only have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv4 forwarding on a VLAN also enables IPv6 hardware forwarding on that VLAN. Future switches may have independent controls per-VLAN for forwarding of IPv4 and IPv6 unicast packets.

broadcast, ignore-broadcast, and fast-directbroadcast each prompt with a warning message when executed while the IP forwarding on the corresponding VLAN is disabled. The hardware and software are NOT programmed until IP forwarding is enabled on the VLAN

fast-direct-broadcast and ignore-broadcast can NOT be enabled simultaneously. These are mutually exclusive.

The broadcast option can be enabled in conjunction with fast-direct-broadcast and ignore-broadcast.

Example

The following command enables forwarding of IP traffic for all VLANs in the current virtual router context with IP addresses:

```
enable ipforwarding
```

The following command enables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
enable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The `ipv4` keyword was added in ExtremeXOS 11.2.

The `ignore-broadcast` and the `fast-direct-broadcast` keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable ip-option loose-source-route

```
enable ip-option loose-source-route
```

Description

Enables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the loose source route IP option (0x83) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

When loose source routing is used, it means that the packet may traverse through other routers that lie in between the designated router path.

With loose source routing enabled, the switch will forward IP packets with the loose source route option enabled, if the switch's IP is in the designated list, or is not in the list.

Example

The following command enables processing of the loose source route IP option:

```
enable ip-option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option record-route

```
enable ip-option record-route
```

Description

Enables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

IP option record-route (IP option 7) means that each router along the path should add its IP address into the options data.

Enabling means that the switch will add itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command enables processing of the record route IP option:

```
enable ip-option record-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option record-timestamp

```
enable ip-option record-timestamp
```

Description

Enables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enables the switch to use the timestamp IP option (0x44). When the switch receives an IP packet with the timestamp option turned on, it will insert the timestamp into the IP options header before forwarding the packet to the destination.

Example

The following command enables processing of the record timestamp IP option:

```
enable ip-option record-timestamp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option strict-source-route

```
enable ip-option strict-source-route
```

Description

Enables processing of the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the strict source route IP option (0x89) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

When strict source routing is used, it means that the packet must use the exact path of routers that lie in the designated router path.

With strict source routing enabled, the switch will forward IP packets with the strict source route option enabled, only if the switch's IP is in the designated list and as long as the next hop in the list is directly attached to one of the router's interfaces.

Example

The following command enables processing of the strict source route IP option:

```
enable ip-option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option router-alert

```
enable ip-option router-alert
```

Description

Enables processing of the router alert IP option in IPv4 packet headers.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables processing of the router alert IP option:

```
enable ip-option router-alert
```

History

This command was first available in EXOS 10.1.

Platform Availability

This command is available on all platforms.

enable iproute compression

```
enable iproute {ipv4} compression {vr <vrname>}
```

Description

Enables IP route compression.

Syntax Description

ipv4	Enable IP route compression algorithm for IPv4 address family. Default is IPv4.
vrname	Virtual router name for which the IP route compression is being enabled. If the virtual router name is not specified, the route compression is enabled for the VR context from which CLI command is being issued.

Default

Disabled.

Usage Guidelines

Enables IP route compression for a specified address family and virtual router. The command applies a compression algorithm on each of the IP prefixes in the routing table.

Example

```
enable iproute compression
```

By default, enables route compression for the IPv4 family for the VR context from which the CLI command is being used.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable iproute sharing

```
enable iproute sharing
```

Description

Enables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost is will be shared.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

IP route sharing allows multiple equal-cost routes to be used concurrently. IP route sharing can be used with static routes or with OSPF routes. In OSPF, this capability is referred to as *equal cost multipath* (ECMP) routing.

Configure static routes and/or OSPF as you would normally. ExtremeXOS supports route sharing across up to 8 static routes or ECMP routes for OSPF. However, on the BlackDiamond 8800 family, SummitStack, and Summit family switches, by default, up to 4 routes are supported. To support up to 8 routes on these switches, use the following command:

```
configure iproute sharing max-gateways <max_gateways>
```

Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic will travel.

Example

The following command enables load sharing for multiple routes:

```
enable iproute sharing
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable irdp

```
enable irdp {vlan <name>}
```

Description

Enables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch will send out IP packets at the intervals specified identifying itself as a default router. IRDP enabled client machines will use this information to determine which gateway address to use for routing data packets to other networks.

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command enables IRDP on VLAN *accounting*:

```
enable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable udp-echo-server

```
enable udp-echo-server {vr <vrid>}{udp-port <port>}
```

Description

Enables UDP echo server support.

Syntax Description

port	Specifies the UDP port.
vr	Specifies the virtual router.

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving end.

Example

The following command enables UDP echo server support:

```
enable udp-echo-server
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

rtlookup

```
rtlookup {<ipaddress> | <ipv6address>}{vr <vrname>}
```

Description

Performs a look-up in the route table to determine the best route to reach an IP address or host.

Syntax Description

ipaddress	Specifies an IPv4 address.
ipv6address	Specifies an IPv6 address.
vrname	Specifies the virtual router to route from.

Default

N/A.

Usage Guidelines

None.

Example

The following command performs a look up in the route table to determine the best way to reach the specified IP address:

```
rtlookup 10.34.25.1
```

History

This command was first available in ExtremeXOS 10.1.

The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show bootprelay

```
show bootprelay
```

Description

Displays the DHCP/BOOTP relay statistics and configuration for the virtual routers.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

None.

Example

The following command displays the DHCP/BOOTP relay statistics for existing virtual routers:

```
show bootprelay
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show iparp

```
show iparp {<ip_addr> | <mac> | vlan <vlan_name> | permanent} {vr
<vr_name>}
```

Description

Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, MAC address, VLAN, or permanent entries.

Syntax Description

ip_addr	Specifies an IP address.
mac	Specifies a MAC address.
vlan_name	Specifies a VLAN name.
permanent	Specifies permanent entries.
vr_name	Specifies a virtual router.

Default

Show all entries, except for proxy entries.

Usage Guidelines

Displays the IP ARP table, including:

- IP address
- MAC address
- Aging timer value
- VLAN name, VLAN ID and port number
- Flags

If you do not specify a virtual router, the command applies to the current virtual router.

The show output displays the following information:

Dynamic entries	The number of dynamic (learned ARP) entries in the table.
Static entries	The number of configured (static ARP) entries in the table.
Pending entries	The number of sent ARP requests that have not yet received a response.
In Request	The number of ARP request packets received (by this virtual router).
In Response	The number of ARP reply packets received (by this virtual router).
Out Request	The number of ARP request packets sent (by this virtual router).
Out Response	The number of ARP reply packets sent (by this virtual router).
Failed requests	The number of failed ARP requests sent (by this virtual router).
Proxy Answered	The number of ARP requests answered by the ARP proxy.

RX Error	The number of incorrect ARP request and reply packets received. The malformed packets include the following errors: incorrect ARP op code, hardware address type is not ethernet, the protocol address is not IP, and similar errors.
Dup IP Addr	IP addresses that have been used by other hosts on the network.
Rejected Count	The number of rejected ARP request packets.
Rejected IP	The source address for the last rejected ARP request. An example reason for an ARP request packet to be rejected is if the source address of the packet is not in the subnet.
Rejected Port	The port on which the last rejected ARP request packet arrived.
Rejected I/F	The VLAN on which the last rejected ARP request packet arrived.
Max ARP entries	Maximum ARP table size for the virtual router (each virtual router has its own ARP table).
Max ARP pending entries	Maximum number of incomplete (pending) ARP entries allowed in the table.
ARP address check	Whether IP ARP checking is enabled or disabled. IP ARP checking verifies if the ARP request's source address is in the receiving interface's subnet.
ARP refresh	Whether ARP refresh is enabled or disabled. ARP refresh is performed when an ARP entry's age is three-fourths of the timeout value.
Timeout	Timeout value for a dynamic (learned) ARP entry.

Example

The following command displays the IP ARP table for the current virtual router:

```
show iparp
```

The following is sample output for the command:

VR	Destination	Mac	Age	Static	VLAN	VID	Port
VR-Default	10.10.10.6	00:04:96:1f:a5:71	8	NO	bluered	4092	1
VR-Default	10.128.32.1	00:01:30:ba:6a:a0	0	NO	Default	4095	
VR-Default	10.128.32.2	00:01:03:1c:ae:b0	5	NO	Default	4095	
VR-Default	10.128.32.4	00:d0:59:17:74:83	3	NO	Default	4095	
VR-Default	10.128.32.5	00:02:a5:c2:5c:dd	0	NO	Default	4095	
VR-Default	10.128.32.6	00:12:3f:1c:f8:fb	5	NO	Default	4095	
VR-Default	10.128.32.7	00:11:11:80:9c:b9	7	NO	Default	4095	
VR-Default	10.128.32.8	00:11:43:53:8e:f1	0	NO	Default	4095	
VR-Default	10.128.32.9	00:02:a5:bf:ac:70	7	NO	Default	4095	
VR-Default	10.128.32.10	00:11:43:44:18:68	10	NO	Default	4095	
VR-Default	10.128.32.11	00:12:3f:1c:e9:f2	0	NO	Default	4095	
VR-Default	10.128.32.12	00:02:a5:bf:af:79	8	NO	Default	4095	
VR-Default	10.128.32.13	00:11:43:40:89:91	0	NO	Default	4095	
VR-Default	10.128.32.16	00:0f:1f:c9:2d:80	2	NO	Default	4095	
VR-Default	10.128.32.17	00:06:5b:b1:6a:91	1	NO	Default	4095	
VR-Default	10.128.32.19	00:11:43:3a:96:1d	10	NO	Default	4095	
VR-Default	10.128.32.20	00:08:02:d5:c5:b7	6	NO	Default	4095	
VR-Default	10.128.32.24	00:12:3f:0a:44:92	14	NO	Default	4095	
VR-Default	10.128.32.26	00:50:04:ad:36:5e	6	NO	Default	4095	
VR-Default	10.128.32.30	00:b0:d0:23:f2:9a	11	NO	Default	4095	
VR-Default	10.128.32.54	00:b0:d0:59:e4:e2	6	NO	Default	4095	
VR-Default	10.128.32.55	00:a0:c9:0c:41:de	3	NO	Default	4095	
VR-Default	10.128.32.59	00:b0:d0:7c:d6:07	14	NO	Default	4095	
VR-Default	10.128.32.99	00:04:96:05:00:03	13	NO	Default	4095	
VR-Default	10.128.32.101	00:04:96:1f:a8:48	0	NO	Default	4095	

```

VR-Default    10.128.32.104    00:30:48:41:ed:45    0    NO    Default    4095
VR-Default    10.128.32.105    00:30:48:41:ed:97    0    NO    Default    4095
VR-Default    10.128.32.106    00:01:30:23:c1:00    0    NO    Default    4095
VR-Default    10.128.32.108    00:04:96:1f:a5:71    0    NO    Default    4095
VR-Default    10.128.32.116    00:04:96:1f:a4:0e    0    NO    Default    4095

Dynamic Entries :          1          Static Entries      :          0
Pending Entries :          0
In Request      :        111          In Response          :          3
Out Request     :        110          Out Response         :        111
Failed Requests :          0
Proxy Answered  :          0
Rx Error        :          0          Dup IP Addr          :          0.0.0.0
Rejected Count  :
Rejected Port   :          Rejected IP          :
          Rejected I/F          :

Max ARP entries :        4096          Max ARP pending entries :        256
ARP address check:  Enabled          ARP refresh           :  Enabled
Timeout         :        20 minutes

```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show iparp proxy

```
show iparp proxy {[<ipNetmask> | <ip_addr> <mask>]} {vr <vr_name>}
```

Description

Displays the proxy ARP table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_address	Specifies an IP address.
mask	Specifies a subnet mask.
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines

If no argument is specified, then all proxy ARP entries are displayed.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays the proxy ARP table:

```
show iparp proxy 10.1.1.5/24
```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show ipconfig

```
show ipconfig {ipv4} {vlan <vlan_name>}
```

Description

Displays configuration information for one or more VLANs.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN(s) information will be displayed. Global IP configuration information includes:

- IP address/netmask/etc.
- IP forwarding information / IP multicast forwarding information
- VLAN name and VLANID
- ICMP configuration (global)
- IRDP configuration (global)
- Maximum ECMP gateways—SummitStack, Summit family, and BlackDiamond 8800 family only

Example

The following command displays configuration information on a VLAN named *accounting*:

```
show ipconfig vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show iproute

```
show iproute {ipv4} {priority | vlan <vlan_name> | permanent | <ip_address>
<netmask> | summary} {multicast | unicast} {vr <vrname>}}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

priority	Specifies a route priority.
vlan_name	Specifies a VLAN name.
permanent	Specifies permanent routing.
ip_address	Specifies an IP address.
netmask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

If a route is active and in use, it is preceded in the display by an “*”. If there are multiple routes to the same destination network, the “*” will indicate which route is the most preferable route. A “c” flag in the Flags column indicates a compressed route resulting from enabling compression using the `enable iproute compression` command. The total number of compressed routes is also shown.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays detailed information about all IP routing:

```
show iproute
```

Output:

```

Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
#be 3.0.0.0/8        111.222.0.5  7    UG-D---um--- feed      0d:19h:52m:49s
#be 4.0.0.0/8        111.222.0.5  5    UG-D---um--- feed      0d:19h:52m:49s
#be 4.0.0.0/9        111.222.0.5  5    UG-D---um--c feed      0d:19h:52m:49s
#be 4.23.84.0/22     111.222.0.5  7    UG-D---um--c feed      0d:19h:52m:49s
#be 4.23.112.0/22    111.222.0.5  7    UG-D---um--c feed      0d:19h:52m:49s
.....
```

```
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
              (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
              (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
```

```

(mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route

```

```

Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
      (m) Multicast, (P) LPM-routing, (R) Modified, (S) Static
      (u) Unicast, (U) Up (c) Compressed

```

Mask distribution:

```

19 routes at length 8          9 routes at length 9
9 routes at length 10         28 routes at length 11

```

Route Origin distribution:

```

7 routes from Direct          184816 routes from EBGp

```

Total number of routes = 184823

Total number of compressed routes = 93274

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

The “c” flag was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show iproute origin

```
show iproute {ipv4} origin [bgp | ebgp | ibgp | direct | static | blackhole
| rip | bootp | icmp | imbgp | ospf | ospf-intra | ospf-inter | ospf-
extern1 | ospf-extern2] {vr <vrname>}
```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

origin	Specifies the origin of the routes to display.

Default

N/A.

Usage Guidelines

Policies for IP routing can be configured based on the route origin. When routes are added to the IP routing table from various sources, the routing policies configured for the origin of the route is applied to the route. After matching on specified characteristics, the characteristics for the route can be modified using the policy. The characteristics that can be matched and modified are dependent on the origin of the route. Policies for IP routing can be dynamically changed. In the case of direct and static route origins, the changes are reflected immediately. In the case of routes that are sourced from other origin, the changes are reflected within 30 seconds.

Example

The following command displays all the BGP routes:

```
show iproute origin bgp
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ipstats

```
show ipstats {ipv4} {vlan <name> | vr <vrname>}
```

Description

Displays IP statistics for the CPU for the switch or for a particular VLAN.

Syntax Description

name	Specifies a VLAN name.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU.

If you do not specify a virtual router, the command applies to the current virtual router.

The fields displayed in the `show ipstats` command are defined in [Table 27](#) through [Table 30](#).

Table 27: Global IP statistics field definitions

Field	Definition
InReceives	Total number of incoming IP packets processed by the CPU.
InUnicast	Total number of unicast IP packets processed by the CPU.
InBcast	Total number of broadcast IP packets processed by the CPU.
InMcast	Total number of multicast IP packets processed by the CPU.
InHdrEr	Total number of packets with an IP Header Error forwarded to the CPU.
Bad vers	Total number of packets with a version other than IP v4 in the IP version field.
Bad chksum	Total number of packets with a bad IP checksum forwarded to the CPU.
Short pkt	IP packets that are too short.
Short hdr	IP packets with a header that is too short.
Bad hdrlen	IP packets with a header length that is less than the length specified.
Bad length	IP packets with a length less than that of the header.
InDelivers	IP packets passed to upper layer protocols.
Bad Proto	IP packets with unknown (not standard) upper layer protocol.
OutRequest	IP packets sent from upper layers to the IP stack.
OutDiscard	IP packets that are discarded due to lack of buffer space or the router interface being down, or broadcast packets with broadcast forwarding disabled.
OutNoRoute	IP packets with no route to the destination.

Table 27: Global IP statistics field definitions (Continued)

Field	Definition
Forwards	ForwardOK and Fwd Err aggregate count.
ForwardOK	Total number of IP packets forwarded correctly.
Fwd Err	Total number of IP packets that cannot be forwarded.
NoFwding	Aggregate number of IP packets not forwarded due to errors.
Redirects	IP packets forwarded on the same network.
No route	Not used.
Bad TTL	IP packets with a bad time-to-live.
Bad MC TTL	IP packets with a bad multicast time-to-live.
Bad IPdest	IP packets with an address that does not comply with the IP v4 standard.
Blackhole	IP packets with a destination that is a blackhole entry.
Output err	Not used. This is the same as Fwd Err.
MartianSrc	IP packets with an invalid source address.

Table 28: Global ICMP statistics field definitions

Field	Definition
OutResp	Echo replies sent from the CPU.
OutError	Redirect from broadcast or multicast source addresses.
InBadcode	Incoming ICMP packets with an invalid CODE value.
InTooshort	Incoming ICMP packets that are too short.
Bad chksum	Incoming ICMP packets with checksum errors.
In Badlen	Incoming ICMP packets with length errors.
echo reply (In/Out):	ICMP “echo reply” packets that are received and transmitted.
destination unreachable (In/Out):	ICMP packets with destination unreachable that are received and transmitted.
port unreachable (In/Out):	ICMP packets with port unreachable that are received and transmitted.
echo (In/Out):	ICMP echo packets that are received and transmitted.

Table 29: Global IGMP statistics field definitions

Field	Definition
Out Query	Number of IGMP query messages sent by the router.
Out Report	Number of reports sent on an active multicast route interface for reserved multicast addresses and for regular IGMP reports forwarded by the query router.
Out Leave	Number of IGMP out leave messages forwarded for IP multicast router interfaces.
In Query	Number of IGMP query messages received.
In Report	Number of IGMP report messages received (mostly from hosts).
In Leave	Number of IGMP leave messages received (mostly from hosts).
In Error	Number of IGMP packets with bad header fields or checksum failures.

Table 30: Router interface statistics field definitions

Field	Definition
Packets IN/OUT	Total number of IP packets received or transmitted on a VLAN router interface.
Octets IN/OUT	Total number of octets received or transmitted on a VLAN router interface.
Mcast packets IN/OUT	Total number of multicast packets received or transmitted on a VLAN router interface.
Bcast packets IN/OUT	Total number of broadcast packets received or transmitted on a VLAN router interface.
Errors IN/OUT	Total number of IP packets with errors received or transmitted on a VLAN router interface.
Discards IN/OUT	Total number of IP packets that cannot travel up to the CPU due to lack of buffer space.
Unknown Protocols IN/OUT	Total number of IP packets with unknown upper layer protocols received by the router interface.

Example

The following command displays IP statistics for the VLAN *accounting*:

```
show ipstats vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The keyword `ipv4` was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show udp-profile

```
show udp-profile {vlan <vlan-name> | {policy} <policy-name>}
```

Description

Displays UDP forwarding profiles.

Syntax Description

vlan-name	Specifies a VLAN.
policy-name	Specifies a UDP forwarding profile.

Default

If no VLAN or policy is specified, all configured profiles will be displayed.

Usage Guidelines

UDP profiles can also be displayed by using the policy manager command `show policy {<policy-name> | detail}`. However, the format of the policy display will be different than this command.

Example

The following command displays all the configured UDP forwarding profiles on the switch:

```
show udp-profile
```

The following is sample output:

```
UDP Profile Name: move_to7
  Number of datagram forwarded: 181
    Dest UDP Port: 67 Fwd to IP Addr: 20.0.0.5
    Dest UDP Port: 67 Fwd to VLAN: to7
  Applied to incoming traffic on VLANS:
    to-mariner
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information check

```
unconfigure bootprelay dhcp-agent information check
```

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to disable the switch from preventing DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To enable this check, use the following command:

```
configure bootprelay dhcp-agent information check
```

Example

The following command disables the DHCP relay agent option check:

```
unconfigure bootprelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information option

```
unconfigure bootprelay dhcp-agent information option
```

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable the DHCP relay agent option (option 82), use the following command:

```
configure bootprelay dhcp-agent information option
```

Example

The following command disables the DHCP relay agent option:

```
unconfigure bootprelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information policy

```
unconfigure bootprelay dhcp-agent information policy
```

Description

Unconfigures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

Replace.

Usage Guidelines

Use this command to unconfigure the policy for the relay agent.

Example

The following command unconfigures the DHCP relay agent option 82 policy:

```
unconfigure bootprelay dhcp-agent information policy
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure icmp

```
unconfigure icmp
```

Description

Resets all ICMP settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all ICMP settings to the default values.

```
unconfigure icmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure iparp

```
unconfigure iparp
```

Description

Resets the following to their default values:

- IP ARP timeout
- max ARP entries
- max ARP pending entries
- ARP checking
- ARP refresh

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets IP ARP timeout to its default value:

```
unconfigure iparp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure irdp

```
unconfigure irdp
```

Description

Resets all router advertisement settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all router advertisement settings to the default values.

```
unconfigure irdp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure udp-profile

```
unconfigure vlan <vlan_name> udp-profile
```

Description

Removes any UDP forwarding profile from a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

None.

Example

The following command removes any UDP forwarding profile from the VLAN *to-sales*:

```
unconfigure vlan to-sales udp-profile
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

The switch provides full Layer 3, IPv6 unicast routing. It exchanges routing information with other routers on the network using either the IPv6 version of Routing Information Protocol (RIPng) or the IPv6 version of Open Shortest Path First (OSPFv3) protocol. The switch dynamically builds and maintains a routing table and determines the best path for each of its routes.

Beginning in release 11.2, ExtremeXOS can provide both IPv4 and IPv6 routing at the same time. Separate routing tables are maintained for the two versions. Most commands that require you to specify an IP address can now accept either an IPv4 or IPv6 address and act accordingly. Additionally, many of the IP configuration, enabling, and display commands have added tokens for IPv4 and IPv6 to clarify the version required. For simplicity, existing commands affect IPv4 by default and require you to specify IPv6, so configurations from an earlier release will still correctly configure an IPv4 network.

This chapter contains commands used with IPv6 unicast routing.

clear neighbor-discovery cache

```
clear neighbor-discovery cache ipv6 {<ipv6address> {vr <vr_name>} | vlan
<vlan_name> | vr <vr_name>}
```

Description

Deletes a dynamic entry from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command clears dynamic entries from the neighbor cache. The `vr` option is used to specify the virtual router on which the operation is performed. When this option is omitted it applies to VR-Default.

When the `ipv6address` or `vlan` options are specified, only the entries with matching IPv6 addresses or that correspond to that VLAN are cleared.

Example

The following command clears all entries from the neighbor cache:

```
clear neighbor-discovery cache
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add

```
configure iproute add <ipv6Netmask> [<ipv6Gateway> | <ipv6ScopedGateway>]
{<metric>} {vr <vrname>} {multicast-only | unicast-only}
```

Description

Adds an IPv6 static route to the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
ipv6Gateway	Specifies a gateway.
ipv6ScopedGateway	Specifies a scoped gateway.
metric	Specifies a cost metric.
vrname	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used. If you do not specify a metric, then the default metric of 1 is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.

Example

The following command adds a static route to the routing table:

```
configure iproute add 2001:db8:0:1111::/64 fe80::1111%default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add blackhole

```
configure iproute add blackhole {ipv6} [<ipv6Netmask>] {vr <vrname>}
{multicast-only | unicast-only}
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
vrname	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry directs packets with a matching specified address prefix to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following command causes packets with a destination address of 2001:db8::3452 to be silently discarded:

```
configure iproute add blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add blackhole ipv6 default

```
configure iproute add blackhole ipv6 default {vr <vrname>} {multicast-only  
| unicast-only}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped.

Syntax Description

vr_name	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is ::.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole ipv6 default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add default

```
configure iproute add default [<ipv6Gateway> | <ipv6ScopedGateway>]
{metric} {vr <vrname>} {multicast-only | unicast-only}
```

Description

Adds a default gateway to the routing table.

Syntax Description

ipv6Gateway	Specifies a VLAN gateway IPv6 address.
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.
ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. Use the `unicast-only` or `multicast-only` options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute delete

```
configure iproute delete <ipv6Netmask> [<ipv6Gateway> |
<ipv6ScopedGateway>] {vr <vrname>}
```

Description

Deletes an IPv6 static route from the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
ipv6Gateway	Specifies a gateway.
ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.

Example

The following command deletes a static address from the routing table:

```
configure iproute delete 2001:db8:0:1111::/64 fe80::1111
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole

```
configure iproute delete blackhole [<ipv6Netmask>] {vr <vrname>}
```

Description

Deletes a blackhole route from the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
vrname	Specifies the virtual router to which the route is added.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry directs packets with a specified destination address to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Example

The following command deletes a blackhole route from the routing table for packets with a destination address of 2001:db8::3452, so the packets are no longer discarded:

```
configure iproute delete blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole ipv6 default

```
configure iproute delete blackhole ipv6 default {vr <vrname>}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

vr_name	Specifies the virtual router to which the route is added.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is ::.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute delete default

```
configure iproute delete default [<ipv6Gateway> | <ipv6ScopedGateway>] {vr
<vrname>}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

ipv6Gateway	Specifies a VLAN gateway IPv6 address.
ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default route from the switch:

```
configure iproute delete default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute ipv6 priority

```
configure iproute ipv6 priority [ripng | blackhole | icmp | static |
ospfv3-intra | ospfv3-inter | ospfv3-as-external | ospfv3-extern1 | ospfv3-
extern2] <priority>
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

ripng	Specifies RIPng.
icmp	Specifies ICMP.
blackhole	Specifies the blackhole route.
static	Specifies static routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-as-external	Specifies OSPFv3 AS External routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
priority	Specifies a priority number in the range of 11 to 65534.

Default

[Table 31](#) lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 31: Relative route priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
RIPg	2400
OSPFv3 AExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.



NOTE

The priority for a blackhole route can not overlap with the priority of any other route origin.

Example

The following command sets the IPv6 route priority for static routing to 1200:

```
configure iproute ipv6 priority static 1200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure neighbor-discovery cache add

```
configure neighbor-discovery cache {vr <vr_name>} add [<ipv6address> |  
<scoped_link_local>] <mac>
```

Description

Adds a static entry to the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
scoped_link_local	Specifies a scoped, link-local address.
mac	Specifies a MAC address.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

This command adds static entries to the neighbor cache.

Example

The following command adds a static entry to the neighbor cache:

```
configure neighbor-discovery cache add fe80::2315%default 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure neighbor-discovery cache delete

```
configure neighbor-discovery cache delete {vr <vr_name>} [<ipv6address> |  
<scoped_link_local>]
```

Description

Deletes a static entry from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
scoped_link_local	Specifies a scoped, link-local address.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

This command deletes static entries from the neighbor cache.

Example

The following command deletes a static entry from the neighbor cache:

```
configure neighbor-discovery cache delete fe80::2315%default
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery add prefix

```
configure vlan <vlan_name> router-discovery {ipv6} add prefix <prefix>
```

Description

Adds a prefix to the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies the prefix to add.

Default

N/A.

Usage Guidelines

This command adds a prefix to the router advertisement messages for the VLAN. Prefixes defined with this command are only included in the router advertisement messages and have no operational impact on VLANs.

To configure the parameters for this prefix, use the following command:

```
configure vlan <vlan_name> router-discovery {ipv6} set prefix <prefix> [autonomous-  
flag <auto_on_off> | onlink-flag <onlink_on_off> | preferred-lifetime <preflife>  
| valid-lifetime <validlife>]
```

Example

The following command adds the prefix 2001:db8:3456::/64 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery add prefix 2001:db8:3456::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery delete prefix

```
configure vlan <vlan_name> router-discovery {ipv6} delete prefix [<prefix>
| all]
```

Description

Deletes prefixes from the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies the prefix to delete.
all	Specifies to delete all prefixes.

Default

N/A.

Usage Guidelines

This command deletes previously defined router advertisement prefixes.

Example

The following command deletes the prefix 2001:db8:3161::/64 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery delete 2001:db8:3161::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery default-lifetime

```
configure vlan <vlan_name> router-discovery {ipv6} default-lifetime  
<defaultlifetime>
```

Description

Configures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
defaultlifetime	Specifies the router lifetime. Range is 0, max-interval to 9000 seconds.

Default

1800 seconds

Usage Guidelines

This command configures the router lifetime value to be included in the router advertisement messages.

The value is specified in seconds and is either 0, or between max-interval and 9000 seconds. A value of 0 indicates that the router is not to be used as a default router.

After a host sends a router solicitation, and receives a valid router advertisement with a non-zero router lifetime, the host must desist from sending additional solicitations on that interface, until an event such as re-initialization takes place.

Example

The following command configures the default-lifetime to be 3600 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery default-lifetime 3600
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery hop-limit

```
configure vlan <vlan_name> router-discovery {ipv6} hop-limit  
<currenthoplimit>
```

Description

Configures the current hop limit value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
currenthoplimit	Specifies the current hop limit. Range is 0 to 255.

Default

64.

Usage Guidelines

Configures the current hop limit value to be included in the router advertisement messages. Hosts then use the current hop limit when they send packets.

A value of 0 means unspecified by this router. The default value is 64. The maximum value is 255.

Example

The following command configures the current hop limit to be 32 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery hop-limit 32
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery link-mtu

```
configure vlan <vlan_name> router-discovery {ipv6} link-mtu <linkmtu>
```

Description

Configures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
linkmtu	Specifies the link MTU. Range is 0 to 9216.

Default

0, meaning that no link MTU information is sent.

Usage Guidelines

This command configures the link MTU placed into the router advertisement messages. Advertisement of the MTU helps ensure use of a consistent MTU by hosts on the VLAN.

The minimum value is 0. The maximum value is 9216. The default value is 0, which means that no link MTU information is included in the router discovery messages.

Example

The following command configures the link MTU to be 5126 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery link-mtu 5126
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery managed-config-flag

```
configure vlan <vlan_name> router-discovery {ipv6} managed-config-flag  
<on_off>
```

Description

Configures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
on_off	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the managed address configuration flag in the router advertisement messages.

A value of on tells hosts to use the administered (stateful) protocol (DHCP) for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration. A value of off tells hosts to use stateless address autoconfiguration. If this command is not entered, the default value is off.

Example

The following command configures the managed address configuration flag to be on for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery managed-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery max-interval

```
configure vlan <vlan_name> router-discovery {ipv6} max-interval  
<maxinterval>
```

Description

Configures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
maxinterval	Specifies the maximum time between advertisements, in seconds. Range is 4 to 1800

Default

600 seconds

Usage Guidelines

This command configures the maximum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

Example

The following command configures the max-interval to be 300 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery max-interval 300
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery min-interval

```
configure vlan <vlan_name> router-discovery {ipv6} min-interval  
<mininterval>
```

Description

Configures the minimum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
mininterval	Specifies the minimum time between advertisements, in seconds. Range is 3 to 1350 (see guidelines).

Default

200 seconds, or max-interval * .33 (see guidelines)

Usage Guidelines

This command configures the minimum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

The minimum value is 3 seconds. The maximum time is $(.75 * \text{max-interval})$ seconds. If you do not explicitly set this value, the min-interval value is reset whenever the max-interval is configured. Min-interval will then be dynamically adjusted to .33 times the max-interval.

Example

The following command configures the min-interval to be 300 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery min-interval 300
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery other-config-flag

```
configure vlan <vlan_name> router-discovery {ipv6} other-config-flag
<on_off>
```

Description

Configures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
on_off	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the other stateful configuration flag in the router advertisement messages.

When set to on, hosts use the administered (stateful) protocol (DHCP) for autoconfiguration of other (non-address) information. If this command is not entered, the default value is off.

Example

The following command configures the other stateful configuration flag to be on for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery other-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery reachable-time

```
configure vlan <vlan_name> router-discovery {ipv6} reachable-time  
<reachabletime>
```

Description

Configures the reachable time value in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
reachabletime	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 3,600,000 (one hour).

Default

0, meaning that no reachable time information is sent.

Usage Guidelines

The reachable time is the time, in milliseconds, that a node assumes a neighbor is reachable after having received a reachability confirmation. A value of 0 means the time is unspecified by this router. The maximum value is 3,600,000 (1 hour). The default value is 0.

Example

The following command configures the reachable time to be 3,600,000 milliseconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery reachable-time 3600000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery retransmit-time

```
configure vlan <vlan_name> router-discovery {ipv6} retransmit-time  
<retransmittime>
```

Description

Configures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
retransmittime	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 4,294,967,295 (approximately 50 days).

Default

0, meaning that no retransmit time information is sent.

Usage Guidelines

This command configures the retransmit time value in the router advertisement messages.

The retransmit time, in milliseconds, is the time between retransmitted neighbor solicitation messages. A value of 0 means the value is unspecified by this router. The maximum value is 4,294,967,295. The default value is 0.

Example

The following command configures the retransmit time to be 604,800,000 milliseconds (one week) for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery retransmit-time 604800000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan router-discovery set prefix

```
configure vlan <vlan_name> router-discovery {ipv6} set prefix <prefix>
[autonomous-flag <auto_on_off> | onlink-flag <onlink_on_off> | preferred-
lifetime <preflife> | valid-lifetime <validlife>]
```

Description

Sets the parameters for a prefix in the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies which prefix's parameters to set.
auto_on_off	Specifies the autonomous flag.
onlink_on_off	Specifies the on link flag.
preflife	Specifies the preferred lifetime in seconds. Maximum value is 4,294,967,295.
validlife	Specifies the valid lifetime in seconds. Maximum value is 4,294,967,295.

Default

The prefix parameter defaults are:

- Valid lifetime—2,592,000 seconds (30 days)
- On-link flag—on
- Preferred lifetime—604,800 seconds (7 days)
- Autonomous flag—on

Usage Guidelines

This command configures the attributes associated with the specified prefix.

The autonomous flag option modifies the autonomous flag of the prefix. The autonomous flag value specifies whether the prefix can be used for autonomous address configuration (on) or not (off).

The onlink flag option modifies the on link flag of the prefix. The on link flag specifies whether the prefix can be used for on link determination (on) or not (off). The default value of the on link flag is on.

The preferred lifetime option modifies the preferred lifetime of a prefix. The preferred lifetime value is the time (from when the packet is sent) that addresses generated from the prefix via stateless address autoconfiguration remain preferred. The maximum value is 4,294,967,295. The default value is 604,800 seconds (7 days).

The valid lifetime option modifies the valid lifetime of a prefix. The valid lifetime value is the time (from when the packet was sent) that the prefix is valid for the purpose of on-link determination. The maximum value is a 4,294,967,295. The default value is 2,592,000 seconds (30 days).

Example

The following command sets the on link parameter of the prefix 3aaa:3161::/64 to off, for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery set prefix 3aaa:3161::/64 onlink-flag off
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure tunnel ipaddress

```
configure tunnel <tunnel_name> ipaddress [{eui64} [<ipv6_address_mask> |
ipv6-link-local]
```

Description

Configures an IPv6 address/prefix on a tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
eui64	Specifies an EUI64 interface identifier for the lower 64 bits of the address.
ipv6_address_mask	Specifies an IPv6 address / IPv6 prefix length.
ipv6-link-local	Specifies the link-local address for a tunnel.

Default

N/A.

Usage Guidelines

This command will configure an IPv6 address/prefix route on the specified tunnel.

6to4 tunnels must follow the standard address requirement. The address must be of the form 2002:<IPv4_source_endpoint>::/16, where <IPv4_source_endpoint> is replaced by the IPv4 source address of the endpoint, in hexadecimal, colon separated form. For example, for a tunnel endpoint located at IPv4 address 10.20.30.40, the tunnel address would be 2002:a14:1e28::/16. In hex, 10 is a, 20 is 14, 30 is 1e and 40 is 28.

6in4 tunnels have no restrictions on their address format or prefix allocations.

Example

The following command configures the 6in4 tunnel link39 with the IPv6 link-local address:

```
configure tunnel link39 ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

create tunnel 6to4

```
create tunnel <tunnel_name> 6to4 source <source-address>
```

Description

Creates an IPv6-to-IPv4 (6to4) tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
source-address	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-to-IPv4 (also known as a 6to4 tunnel), and add it to the system. A maximum of 1 6to4 tunnel can be configured on any particular virtual router. The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or virtual routers. The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel that uses it still exists.

Example

The following command creates the 6to4 tunnel link35 with source address 192.168.10.1:

```
create tunnel link35 6to4 source 192.168.10.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

create tunnel ipv6-in-ipv4

```
create tunnel <tunnel_name> ipv6-in-ipv4 destination <destination-address>  
source <source-address>
```

Description

Creates an IPv6-in-IPv4 (6in4) tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
source-address	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-in-IPv4 (otherwise known as a configured tunnel or a 6in4 tunnel) and add it to the system. A maximum of 255 tunnels (including one 6to4 tunnel) can be configured on the system. The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or virtual routers. The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel is still exists that uses it.

Example

The following command creates the 6in4 tunnel link39 with destination address 10.10.10.10 and source address 192.168.10.15:

```
create tunnel link39 ipv6-in-ipv4 destination 10.10.10.10 source 192.168.10.15
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

delete tunnel

```
delete tunnel <tunnel_name>
```

Description

Deletes an IPv6 tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

This command will destroy a previously created tunnel. The command acts on either a 6to4 or a 6in4 tunnel. Because the interface will be removed, all routes pointing out that interface will also be purged from the system at the time.

Example

The following command deletes the tunnel link39:

```
delete tunnel link39
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ipforwarding ipv6

```
disable ipforwarding ipv6 {vlan <vlan_name> | tunnel <tunnel_name> | vr  
<vr_name>}
```

Description

Disables routing for one or all interfaces. If no argument is provided, disables routing for all interfaces on the current virtual router.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnel_name	Specifies an IPv6 tunnel.
vr_name	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IP forwarding is disabled by default.

Extreme Networks switches have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.

Example

The following command disables forwarding of IPv6 traffic for a VLAN named *accounting*:

```
disable ipforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable iproute ipv6 compression

```
disable iproute ipv6 compression {vr <vrname>}
```

Description

This command disables IPv6 route compression.

Syntax Description

vr <vrname>	Specifies the virtual router to route from.

Default

The defaults are:

- Vname—current CLI context VR
- Address family—Ipv4

By default, IP route compression is disabled for all address families and VRs.

Usage Guidelines

This command disables IP route compression for the IPv6 address family and VR. This command decompresses previously compressed prefixes in the IPv6 prefix database.

Example

The following example disables IP route compression for the IPv6 address family and the VR of the current CLI context.

```
disable iprouteipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable router-discovery

```
disable router-discovery {ipv6} vlan <vlan_name>
```

Description

Disables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

None

Example

The following command disables router discovery for the VLAN *top_floor*:

```
disable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ipforwarding ipv6

```
enable ipforwarding ipv6 {vlan <vlan_name> | tunnel <tunnel-name> | vr  
<vr_name>}
```

Description

Enables IP routing VLANs. If no argument is provided, enables IP routing for all VLANs and tunnels that have been configured with an IPv6 address on the current virtual router.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnel_name	Specifies an IPv6 tunnel.
vr_name	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IP forwarding is disabled by default.

Extreme Networks switches have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.

Example

The following command enables forwarding of IPv6 traffic for all VLANs in the current virtual router context with IPv6 addresses:

```
enable ipforwarding ipv6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable iproute ipv6 compression

```
enable iproute ipv6 compression {vr <vrname>}
```

Description

This command enables IPv6 route compression.

Syntax Description

vr <vrname>	Specifies the virtual router to route from.

Default

The defaults are:

- Vname—current CLI context VR
- Address family—Ipv4

Usage Guidelines

This command enables IP route compression for the IPv6 address family and VR. This command applies a compression algorithm to each IPv6 prefix in the IPv6 prefix database.

Example

The following example enables IP route compression for the IPv6 address family and the VR of the current CLI context.

```
enable iprouteipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable router-discovery

```
enable router-discovery {ipv6} vlan <vlan_name>
```

Description

Enables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command is only valid when the specified VLAN has an IPv6 address associated with it. After IPv6 Router Discovery is enabled on a VLAN, router advertisement messages are regularly sent on all ports associated with the VLAN.

Example

The following command enables router discovery for the VLAN *top_floor*:

```
enable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

rtlookup

```
rtlookup [<ipaddress> | <ipv6address>]{vr <vrname>}
```

Description

Performs a look-up in the route table to determine the best route to reach an IP address or host.

Syntax Description

ipaddress	Specifies an IPv4 address.
ipv6address	Specifies an IPv6 address.
vrname	Specifies the virtual router to route from.

Default

N/A.

Usage Guidelines

None.

Example

The following command performs a look up in the route table to determine the best way to reach the specified IP address:

```
rtlookup 3aaa:5643::ef80:2525:1023:5213
```

History

This command was first available in ExtremeXOS 10.1.

The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ipconfig ipv6

```
show ipconfig ipv6 {vlan <vlan_name> | tunnel <tunnelname>}
```

Description

Displays configuration information for one or more interfaces.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnelname	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

If no interface is specified, then global IP configuration is displayed. Otherwise, specific interface(s) will be displayed. Global IP configuration information includes:

- IP address/netmask/etc.
- IP forwarding information / IP multicast forwarding information

Example

The following command displays configuration information on a VLAN named *accounting*:

```
show ipconfig ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show iproute ipv6

```
show iproute ipv6 {priority | vlan <vlan_name> | tunnel <tunnel-name> |
<ipv6Netmask> | summary {multicast | unicast}} {vr <vrname>}}
```

Description

Displays the contents of the IPv6 routing table.

Syntax Description

priority	Specifies to display route priority.
vlan_name	Specifies a VLAN name.
tunnel	Specifies a tunnel name.
ipv6Netmask	Specifies an IPv6 address/prefix length.
summary	Specifies summary information

Default

N/A.

Usage Guidelines

If a route is active and in use, it is preceded in the display by an “*”. If there are multiple routes to the same destination network, the “*” will indicate which route is the most preferable route.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays detailed information about all IPv6 routing:

```
show iproute ipv6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show iproute ipv6 origin

```
show iproute ipv6 origin [direct | static | blackhole | ripng | ospfv3 |  
ospfv3-intra | ospfv3-inter | ospfv3-extern1 | ospfv3-extern2] {vr <vrname>}
```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

origin	Specifies a display of the route origin.

Default

N/A.

Usage Guidelines

Policies for IP routing can be configured based on the route origin. When routes are added to the IP routing table from various sources, the routing policies configured for the origin of the route is applied to the route. After matching on specified characteristics, the characteristics for the route can be modified using the policy. The characteristics that can be matched and modified are dependent on the origin of the route. Policies for IP routing can be dynamically changed. In the case of direct and static route origins, the changes are reflected immediately. In the case of routes that are sourced from other origin, the changes are reflected within 30 seconds.

Example

The following command displays the RIPng routes:

```
show iproute ipv6 origin ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ipstats ipv6

```
show ipstats ipv6 {vlan <name> | tunnel <tunnelname> | vr <vrname>}
```

Description

Displays IPv6 statistics for the CPU for the switch or for a particular VLAN.

Syntax Description

name	Specifies a VLAN name.
tunnelname	Specifies a tunnel name.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays IPv6 statistics for the VLAN *accounting*:

```
show ipstats ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show neighbor-discovery cache ipv6

```
show neighbor-discovery cache ipv6 {[[<ipv6_addr> | <mac> | permanent] {vr
<vr_name>}} | vlan <vlan_name> | vr <vr_name>}
```

Description

This command displays all the entries from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
mac	Specifies a MAC address.
permanent	Specifies static entries.
ipv6_addr	Specifies an IPv6 address.
vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command displays the entries present in the neighbor cache.

The entries displayed can be filtered by IPv6 address, MAC address, or by VLAN. The `permanent` keyword filters the output to display static entries.

The `vr_name` indicates the virtual router on which the operation is performed. In its absence, the operation applies to VR-Default.

Example

The following command shows all entries from the neighbor cache:

```
show neighbor-discovery cache ipv6
```

The following is sample output:

VR	Destination	Mac	Age	Static	VLAN	VID	Port
VR-Default	3ffe:100::7	00:01:30:00:6b:00	0	NO	gtag100	100	1:2
VR-Default	3ffe:100::99	00:01:02:33:33:33	0	YES	gtag100	100	
VR-Default	3ffe:99::99	00:01:02:01:01:01	0	YES	gtag99	99	

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show router-discovery

```
show router-discovery {ipv6} {vlan <vlan_name>}
```

Description

Displays the router discovery settings.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

If no VLAN is specified, the settings are displayed for all IPv6 configured VLANs.

Example

The following command displays router discovery settings for the VLAN *top_floor*:

```
show router-discovery vlan top_floor
```

The following is sample output:

```
Router Advertisements disabled on v1
  Minimum/Maximum Interval: 200 / 600
  Managed / Other Info Flags: Off / Off
  Link MTU: 0
  Reachable Time: 0
  Retrans Timer: 0
  Current Hop Limit: 64
  Default Lifetime: 1800
  Number of Prefixes: 1, Prefix List:
                                Valid
                                Prefix Lifetime
11::1/64                        2592000
                                Auto Preferred
                                On 604800 On
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show tunnel

```
show tunnel
```

Description

Displays system tunnel information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays system tunnel information.

Example

The following command displays system tunnel information:

```
show tunnel
```

The following is sample output:

Name	Type	Flags
tunfour	6in4 10.20.30.40 => 10.10.10.10	U
tunfive2	6to4 10.20.30.40 => *.*.*.*	D
Total tunnels: 2		
Flags: (U) Up / (D) Down		

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure tunnel

```
unconfigure tunnel <tunnel_name> ipaddress <ipv6_address_mask>
```

Description

Unconfigures an IPv6 address/prefix route from a tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
ipv6_address_mask	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

This command will unconfigure an IPv6 address/prefix route from the specified tunnel.

Example

The following command unconfigures the 6in4 tunnel link39 with the address 3aaa::1111/64

```
unconfigure tunnel link39 3aaa::1111/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

This chapter describes commands used for the interior gateway protocol RIP.

Routing Information Protocol (RIP) is an Interior Gateway Protocol (IGP) first used in computer routing in the Advanced Research Projects Agency Network (ARPAnet) as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address of the destination network
- Metric (hop count) to the destination network
- IP address of the next router
- Timer that tracks the amount of time since the entry was last updated

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

A new version of RIP, called RIP version 2 (RIPv2), expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs)
- Next-hop addresses
- Support for next-hop addresses allows for optimization of routes in certain environments
- Multicasting

If you are using RIP with supernetting/Classless Inter-Domain Routing (CIDR), you must use RIPv2 only, and RIP route aggregation must be turned off.

clear rip counters

```
clear rip counters
```

Description

Clears the RIP counters (statistics).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIP statistics counters:

```
clear rip counters
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip add vlan

```
configure rip add vlan [<vlan-name> | all]
```

Description

Configures RIP on an IP interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled, the parameters are not reset to default automatically.

Example

The following command configures RIP on the VLAN *finance*:

```
configure rip add finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip delete vlan

```
configure rip delete vlan [<vlan-name> | all]
```

Description

Disables RIP on an IP interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled by this command, the parameters are not reset to default automatically.

Example

The following command deletes RIP on a VLAN named *finance*:

```
configure rip delete finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip garbagetime

```
configure rip garbagetime {<seconds>}
```

Description

Configures the RIP garbage time.

Syntax Description

seconds	Specifies a time in seconds.

Default

120 seconds.

Usage Guidelines

None.

Example

The following command configures the RIP garbage time to have a 60-second delay:

```
configure rip garbagetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip import-policy

```
configure rip import-policy [<policy-name> | none]
```

Description

Associates or removes the policy applied to RIP routes added to the system routing table.

Syntax Description

policy-name	Specifies the policy to apply.

Default

No policy.

Usage Guidelines

Use this command to associate a policy with the RIP routes installed into the system table. Use the `none` option to remove the policy association.

Example

The following example applies the policy *campuseast* to RIP routes:

```
configure rip import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip routetimeout

```
configure rip routetimeout <seconds>
```

Description

Configures the route timeout period.

Syntax Description

seconds	Specifies a time in seconds.

Default

180 seconds.

Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

Example

The following example sets the route timeout period to 120 seconds:

```
configure rip routetimeout 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip vlan rxmode

```
configure rip [vlan <vlan-name> | all] rxmode [none | v1only | v2only | any]
```

Description

Changes the RIP receive mode for one or all VLANs.

Syntax Description

none	Specifies to drop all received RIP packets.
v1only	Specifies to accept only RIP version 1 format packets.
v2only	Specifies to accept only RIP version 2 format packets.
any	Specifies to accept RIP version 1 and RIP version 2 packets.
vlan-name	Specifies to apply settings to specific VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the receive mode for the VLAN *finance* to accept only RIP version 1 format packets:

```
configure rip finance rxmode v1only
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip vlan txmode

```
configure rip [vlan <vlan-name> | all] txmode [none | v1only | v1comp | v2only]
```

Description

Changes the RIP transmission mode for one or all VLANs.

Syntax Description

none	Specifies to not transmit any packets on this interface.
v1only	Specifies to transmit RIP version 1 format packets to the broadcast address.
v1comp	Specifies to transmit RIP version 2 format packets to the broadcast address.
v2only	Specifies to transmit RIP version 2 format packets to the RIP multicast address.
vlan-name	Specifies to apply settings to a specific VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the transmit mode for the VLAN *finance* to transmit version 2 format packets to the broadcast address:

```
configure rip finance txmode v1comp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip updatetime

```
configure rip updatetime <seconds>
```

Description

Specifies the time interval in seconds within which RIP sends update packets.

Syntax Description

seconds	Specifies a time in seconds. The range is 10 to 180.

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). The timer granularity is 10 seconds. Timer minimum is 10 seconds and maximum is 180 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure rip updatetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip vlan cost

```
configure rip vlan [<vlan-name> | all] cost <cost>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
cost	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN *finance* to a metric of 3:

```
configure rip vlan finance cost 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip vlan route-policy

```
configure rip vlan [<vlan-name> | all] route-policy [in | out] [<policy-name> | none]
```

Description

Configures RIP to ignore certain routes received from its neighbor, or to suppress certain routes when performing route advertisements.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
policy-name	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIP routes are accepted as valid routes. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIP routes are advertised on the VLAN.

Example

The following command configures the VLAN *backbone* to accept selected routes from the policy *nosales*:

```
configure rip vlan backbone route-policy in nosales
```

The following command uses the policy *nosales* to determine which RIP routes are advertised into the VLAN *backbone*:

```
configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure rip vlan trusted-gateway

```
configure rip vlan [<vlan-name> | all] trusted-gateway [<policy-name> | none]
```

Description

Configures a trusted neighbor policy to determine trusted RIP router neighbors for the VLAN on the switch running RIP.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
policy-name	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

N/A.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIP control packets from trusted neighbors will be processed.

Example

The following command configures RIP to use the policy *nointernet* to determine from which RIP neighbor to receive (or reject) the routes to the VLAN *backbone*:

```
configure rip vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip

```
disable rip
```

Description

Disables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

Example

The following command disables RIP for the whole router:

```
disable rip
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip aggregation

```
disable rip aggregation
```

Description

Disables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) router.

Syntax Description

This command has no arguments or variables.

Default

RIP aggregation is disabled by default.

Usage Guidelines

The disable RIP aggregation command disables the RIP aggregation of subnet information on a switch configured to send RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Within a class boundary, no routes are aggregated.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command disables RIP aggregation on the interface:

```
disable rip aggregation
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip export

```
disable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 |
ospf-extern2 | ospf-inter | ospf-intra | static]
```

Description

Disables RIP from redistributing routes from other routing protocols.

Syntax Description

static	Specifies static routes.
bgp	Specifies BGP routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies external BGP routes.
i-bgp	Specifies internal BGP routes.
ospf	Specifies all OSPF routes.
ospf-intra	Specifies OSPF-intra area routes.
ospf-inter	Specifies OSPF-inter area routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.

Default

Disabled.

Usage Guidelines

This command disables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain.

Example

The following command disables RIP from redistributing any routes learned from OSPF:

```
disable rip export ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip originate-default

```
disable rip originate-default
```

Description

Disables the advertisement of a default route.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIP if no other default route is advertised:

```
disable rip originate-default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip poisonreverse

```
disable rip poisonreverse
```

Description

Disables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIP:

```
disable rip poisonreverse
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip splithorizon

```
disable rip splithorizon
```

Description

Disables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIP:

```
disable rip splithorizon
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip triggerupdates

```
disable rip triggerupdates
```

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command disables the trigger update mechanism:

```
disable rip triggerupdate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable rip use-ip-router-alert

```
disable rip use-ip-router-alert
```

Description

Disables router alert IP option in outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the RIP router alert IP option:

```
disable rip use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip

```
enable rip
```

Description

Enables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

Example

The following command enables RIP for the whole router:

```
enable rip
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip aggregation

```
enable rip aggregation
```

Description

Enables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) interface.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The enable (disable) rip aggregation command enables (disables) the RIP aggregation of subnet information on an interface configured to send RIPv1 or RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Subnet routes are aggregated to the nearest class network route when crossing a class boundary.
- Within a class boundary, no routes are aggregated.
- If aggregation is enabled, the behavior is the same as in RIPv1.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command enables RIP aggregation on the interface:

```
enable rip aggregation
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip export

```
enable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 |
ospf-extern2 | ospf-inter | ospf-intra | static] [cost <number> {tag
<number>} | policy <policy-name>]
```

Description

Enables RIP to redistribute routes from other routing functions.

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
ospf	Specifies all OSPF routes.
ospf-intra	Specifies OSPF-intra area routes.
ospf-inter	Specifies OSPF-inter area routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
static	Specifies static routes.
cost <number>	Specifies the <code>cost</code> metric, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
tag <number>	Specifies a tag number.
<policy-name>	Specifies a policy.

Default

Disabled.

Usage Guidelines

This command enables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain. You can choose which types of OSPF routes are injected, or you can simply choose `ospf`, which will inject all learned OSPF routes regardless of type.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. If the cost metric is set to 0, the cost is inserted from the route. For example, with BGP, the cost could be the MED or the length of the BGP path. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes.

Example

The following command enables RIP to redistribute routes from all OSPF routes:

```
enable rip export ospf cost 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip originate-default cost

```
enable rip originate-default {always} cost <number> {tag<number>}
```

Description

Configures a default route to be advertised by RIP.

Syntax Description

always	Specifies to always advertise the default route.
cost <number>	Specifies a cost metric. The range is 1 - 15.
tag <number>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If `always` is specified, RIP always advertises the default route to its neighbors. If `always` is not specified, RIP advertises a default route only if a reachable default route is in the system route table.

The default route advertisement is filtered using the out policy.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIP if there is a default route in the system routing table:

```
enable rip originate-default cost 7
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip poisonreverse

```
enable rip poisonreverse
```

Description

Enables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command enables the split horizon with poison reverse algorithm for RIP:

```
enable rip poisonreverse
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip splithorizon

```
enable rip splithorizon
```

Description

Enables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIP:

```
enable rip splithorizon
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip triggerupdates

```
enable rip triggerupdates
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command enables the trigger update mechanism:

```
enable rip triggerupdate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable rip use-ip-router-alert

```
enable rip use-ip-router-alert
```

Description

Enables the router alert IP option in the outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the RIP router alert IP option:

```
enable rip use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show rip

```
show rip
```

Description

Displays RIP specific configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific configuration:

```
show rip
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show rip interface

```
show rip interface {detail}
```

Description

Displays RIP-specific configuration and statistics for all VLANs.

Syntax Description

detail	Specifies detailed display.

Default

Show summary output for all interfaces.

Usage Guidelines

Summary includes the following information per interface:

- VLAN name
- IP address and mask
- interface status
- packets transmitted
- packets received
- number of triggered updates
- cost

Detail includes the following per interface:

- VLAN name
- IP address and mask
- tx mode
- rx mode
- cost
- peer information (for each peer)
 - age
 - version
 - received packets
 - received updates
 - received bad packets
 - received bad routes
- in policy
- out policy

- trusted gateway policy
- packets transmitted
- sent triggered updates
- packets received
- bad packets received
- bad routes received

Example

The following command displays the RIP configuration for all VLANs:

```
show rip interface
```

The following command displays RIP-specific statistics for all VLANs:

```
show rip interface detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show rip interface vlan

```
show rip interface vlan <vlan-name>
```

Description

Displays RIP specific statistics and configuration for a VLAN in detail.

Syntax Description

vlan-name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific statistics for the VLAN *accounting*:

```
show rip interface accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show rip memory

```
show rip memory {detail | <memoryType>}
```

Description

Displays RIP specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific memory for all types:

```
show rip memory detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show rip routes

```
show rip routes {detail} {network <ripNetworkPrefix>}
```

Description

Displays routes advertised by RIP.

Syntax Description

detail	Displays all available information from the RIP routing table.
ripNetworkPrefix	Specifies the route prefix for the routes to show.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIP, including routes exported from the system routing table and originated by other protocols, for example BGP.

Example

The following command displays a summary of RIP specific routes for the networks 10.0.0.0/8:

```
show rip routes network 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure rip

```
unconfigure rip {vlan <vlan-name> | all}
```

Description

Resets all RIP parameters to the default for all VLANs or for the specified VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.

Default

All.

Usage Guidelines

Does not change the enable/disable state of the RIP settings.

Example

The following command resets the RIP configuration to the default for the VLAN *finance*:

```
unconfigure rip finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

31 RIPng Commands

This chapter describes commands used for the IPv6 interior gateway protocol RIPng.

To determine the best path to a distant network, a router using RIPng always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIPng contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address and prefix length of the destination network
- Metric (hop count) to the destination network
- IP address of the next hop router, if the destination is not directly connected
- Interface for the next hop
- Timer that tracks the amount of time since the entry was last updated
- A flag that indicates if the entry is a new one since the last update
- The source of the route, for example, static, RIPng, OSPFv3, etc.

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

clear ripng counters

```
clear ripng counters {vlan <vlan-name> | tunnel <tunnel-name>}
```

Description

Clears the RIPng global or interface-specific counters (statistics).

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIPng statistics counters:

```
clear ripng counters
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng add

```
configure ripng add [vlan <vlan-name> | tunnel <tunnel-name>
| [vlan | tunnel] all]
```

Description

Configures RIPng on an IP interface.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

For RIPng to be active on the interface, it must also be globally enabled using the command [enable ripng](#). If the keyword `all` is specified, all IPv6 configured VLANs or tunnels will be configured for RIPng.

Example

The following command configures RIPng on the VLAN *finance*:

```
configure ripng add finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng delete

```
configure ripng delete [vlan <vlan-name> | tunnel <tunnel-name>
| [vlan | tunnel] all]
```

Description

Removes an interface from RIPng routing.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

This command removes an interface from RIPng routing. However, the RIPng-specific interface configuration will be preserved, even if RIPng is unconfigured on the interface. The interface configuration information is removed only when the IPv6 interface itself gets deleted by, for example, by unconfiguring all the IPv6 addresses on the interface.

Example

The following command removes the VLAN *finance* from RIPng routing:

```
configure ripng delete finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng garbage-time

```
configure ripng garbage-time {<seconds>}
```

Description

Configures the RIPng garbage time.

Syntax Description

seconds	Specifies a time in seconds. Range is 10 to 2400 seconds.

Default

120 seconds.

Usage Guidelines

This command configures the time interval after which a route in the RIPng routing database that has expired will be removed. The value is rounded off to nearest multiple of 10.

Example

The following command configures the RIPng garbage time to have a 60-second delay:

```
configure ripng garbage-time 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng import-policy

```
configure ripng import-policy [<policy-name> | none]
```

Description

Associates or removes the policy applied to RIPng routes added to the system routing table.

Syntax Description

policy-name	Specifies the policy to apply.

Default

No policy.

Usage Guidelines

Use this command to configure the policy to be applied to RIPng routes installed into the system routing table from the RIPng routing process. This policy can be used to modify parameters associated with routes installed into the routing table.

Use the `none` option to remove the policy association.

The following is a sample policy file that can be used with RIPng. It changes the metric to 12 for any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  If match any{
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    cost 12;
  }
}
```

Example

The following example applies the policy *campuseast* to RIPng routes:

```
configure ripng import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng routetimeout

```
configure ripng routetimeout <seconds>
```

Description

Configures the route timeout period for RIPng.

Syntax Description

seconds	Specifies a time in seconds. Range is 10 to 3600.

Default

180 seconds.

Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

The configured value is rounded off to the nearest multiple of 10.

Example

The following example sets the route timeout period to 120 seconds:

```
configure ripng routetimeout 120
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng updatetime

```
configure ripng updatetime <seconds>
```

Description

Specifies the time interval in seconds within which RIPng sends update packets.

Syntax Description

seconds	Specifies a time in seconds. The range is 10 to 3600.

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). The timer granularity is 10 seconds. Timer minimum is 10 second and maximum is 3600 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure ripng updatetime 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng cost

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] cost <metric>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
metric	Specifies a cost metric. Range is 1 to 15.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN *finance* to a metric of 3:

```
configure ripng vlan finance cost 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng route-policy

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] route-policy [in
| out] [<policy-name> | none]
```

Description

Configures RIPng to ignore or modify certain routes received from its neighbors, or to suppress certain routes when performing route advertisements.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
policy-name	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIPng routes are accepted as valid routes from RIPng neighbors. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIPng routes are advertised to other RIPng neighbors.

The following is a sample policy file that could be used with RIPng. It will drop any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  If match any{
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    deny;
  }
}
```

Example

The following command configures the VLAN *backbone* to accept routes from its neighbor as specified by the policy *nosales*:

```
configure ripng vlan backbone route-policy in nosales
```

The following command uses the policy *nosales* to determine which RIP routes are advertised into the VLAN *backbone*:

```
configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ripng trusted-gateway

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] trusted-gateway
[<policy-name> | none]
```

Description

Configures a trusted neighbor policy to determine trusted RIPng router neighbors for the interfaces on the switch running RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
policy-name	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

None. Control packets from all of the neighbors are processed.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIPng control packets from trusted neighbors will be processed.

The following policy designates neighbors from the fe80:202:b3ff:fe4a:6ada:: /64 subnet and the neighbor at fe80:203::b3ff:fe4a:6ada as trusted gateways:

```
entry filter_gateways {
    If match any{
        nlri fe80:202:b3ff:fe4a:6ada:: /64;
        nlri fe80:203::b3ff:fe4a:6ada:: /64;
    }
    then {
        permit;
    }
}
```

Example

The following command configures RIPng to use the policy *nointernet* to determine from which RIPng neighbor to receive (or reject) the routes to the VLAN *backbone*:

```
configure ripng vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng

```
disable ripng
```

Description

Disables RIPng for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables RIPng for the whole router:

```
disable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng export

```
disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 |
ospfv3-inter | ospfv3-intra | static]
```

Description

Disables RIPng from redistributing routes from other routing protocols.

Syntax Description

static	Specifies user configured static routes.
direct	Specifies directly reachable subnets from the router (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.

Default

Disabled.

Usage Guidelines

This command disables the exporting of static, direct, and OSPF-learned routes from the switch routing table into the RIPng domain.

Example

The following command disables RIPng from redistributing any routes learned from OSPFv3:

```
disable ripng export ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng originate-default

```
disable ripng originate-default
```

Description

Disables the advertisement of a default route to the neighbors.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIPng if no other default route is advertised:

```
disable ripng originate-default
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng poisonreverse

```
disable ripng poisonreverse {vlan <vlan-name> | tunnel <tunnel_name> |  
[vlan | tunnel] all}
```

Description

Disables poison reverse algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIPng:

```
disable ripng poisonreverse
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng splithorizon

```
disable ripng splithorizon {vlan <vlan-name> | tunnel <tunnel_name> | [vlan  
| tunnel] all}
```

Description

Disables the split horizon algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIPng:

```
disable rip splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable ripng triggerupdate

```
disable ripng triggerupdate {vlan <vlan-name> | tunnel <tunnel_name> |  
[vlan | tunnel] all}
```

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPng-related traffic.

When this feature is disabled, any metric change on the interface, or an interface going down will not be communicated until the next periodic update. To configure how often periodic updates are sent, use the following command:

```
configure ripng updatetime
```

Example

The following command disables the trigger update mechanism:

```
disable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng

```
enable ripng
```

Description

Enables RIPng for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Although RIPng is useful in small networks, it has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

For larger networks, consider OSPFv3 as an alternative IGP.

Example

The following command enables RIPng for the whole router:

```
enable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng export

```
enable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 |
ospfv3-inter | ospfv3-intra | static] [cost <number> {tag <number>} |
policy <policy-name>]
```

Description

Enables RIPng to redistribute routes from other routing functions.

Syntax Description

direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
static	Specifies static routes.
cost <number>	Specifies the <code>cost</code> metric, from 0-15. If set to 0, RIPng uses the route metric obtained from the route origin.
tag <number>	Specifies a tag number.
<policy-name>	Specifies a policy.

Default

Disabled. However, direct routes will always be advertised for all the interfaces where RIPng is enabled. For those interfaces where RIPng is not enabled, the corresponding direct route could be redistributed if direct route export is enabled through this command.

Default tag is 0.

Usage Guidelines

This command enables the exporting of static, direct, and OSPFv3-learned routes from the routing table into the RIPng domain. You can choose which types of OSPFv3 routes are injected, or you can simply choose `ospfv3`, which will inject all learned OSPFv3 routes regardless of type.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. If the cost metric is set to 0, the cost is inserted from the route table. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes.

The following is sample policy file which modifies the cost of redistributed routes from OSPFv3 and statically configured routes:

```
entry filter_rt {  
    If match any {  
        Route-origin ospfv3;  
        Route-origin static;  
    }  
    then {  
        cost 10;  
    }  
}
```

Example

The following command enables RIPng to redistribute routes from all OSPFv3 routes:

```
enable ripng export ospfv3 cost 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng originate-default

```
enable ripng originate-default {always} cost <metric> {tag <number>}
```

Description

Configures a default route to be advertised by RIPng.

Syntax Description

always	Specifies to advertise the default route in addition to learned default route.
cost <metric>	Specifies a cost metric. The range is 1 - 15.
tag <number>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If `always` is specified, RIPng always advertises the default route to its neighbors. If `always` is not specified, RIPng advertises a default route only if a reachable default route is in the system route table (the route is learned from other neighbors).

The default route advertisement is filtered using the out policy. Use the command, `configure ripng route-policy`, to specify the out policy.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIPng if there is a default route in the system routing table:

```
enable ripng originate-default cost 7
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng poisonreverse

```
enable ripng poisonreverse {vlan <vlan-name> | tunnel <tunnel_name> | [vlan  
| tunnel] all}
```

Description

Enables the split horizon with poison reverse algorithm for RIPng on specified interfaces.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Used with split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

If both split horizon and poison reverse are enabled, poison reverse takes precedence.

Example

The following command enables split horizon with poison reverse for RIPng on all IPv6 interfaces in the virtual router:

```
enable ripng poisonreverse
```

The following command enables split horizon with poison reverse for all the IPv6 configured VLANs in the virtual router:

```
enable ripng poisonreverse vlan all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng splithorizon

```
enable ripng splithorizon {vlan <vlan-name> | tunnel <tunnel_name> | [vlan  
| tunnel] all}
```

Description

Enables the split horizon algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIPng on all IPv6 configured interfaces:

```
enable ripng splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ripng triggerupdates

```
enable ripng triggerupdates {vlan <vlan-name> | tunnel <tunnel_name> |  
[vlan | tunnel] all}
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPng-related traffic.

Example

The following command enables the trigger update mechanism on all IPv6 configured interfaces:

```
enable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ripng

```
show ripng
```

Description

Displays RIPng global configuration and runtime information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIPng global configuration and runtime information:

```
show ripng
```

The following is sample output from this command:

```
Ripng Routing :Enabled
Triggered Updates: Enabled           Aggregation           : Disabled
Update Interval  : 30                 Route Timeout          : 180
Garbage Timeout  : 120                 Hold Down Time         : 0
Originate Default: Disabled
Sys Import-Policy: None
Redistribute:
```

Protocol	Status	Cost	Tag	Policy
Direct	Enabled	0	0	none
Static	Disabled	0	0	none
OSPFIntra	Disabled	0	0	none
OSPFInter	Disabled	0	0	none
OSPFExt1	Disabled	0	0	none
OSPFExt2	Disabled	0	0	none

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ripng interface

```
show ripng interface {detail | vlan <vlan-name> | tunnel <tunnel-name>}
```

Description

Displays RIPng-specific configuration and statistics for the specified interface.

Syntax Description

detail	Specifies detailed display.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Show summary output for all interfaces.

Usage Guidelines

Displays the RIPng interface configuration and runtime information. If no interface is specified, only the summary data for all the configured interfaces is displayed. If an interface is specified, only the data for that interface is displayed in detail. If the keyword `detail` is specified, detailed data for all interfaces is displayed.

Example

The following command displays the RIPng configuration summary for all interfaces:

```
show ripng interface
```

The following is sample output from this command:

VLAN	IP Address	Flags	Sent Packets	Rcvd Packets	Triggered Updates	Cost
v1	22cc::3	/64 rif-pst	106349	106349	3	15
v2	22bb::1	/64 rif-pst	106349	106095	3	1
v3	2abc::1	/120 rif-pst	106351	0	4	1
v4	3ffe::1	/64 rif-pst	106349	139124	3	1

Flags: (f) Interface Forwarding Enabled, (i) Interface RIPng Enabled
 (n) Multinetted Interface, (r) Router RIPng Enabled
 (p) Poison Reverse Enabled, (s) Split Horizon Enabled
 (t) Triggerred Update Enabled.

The following command displays RIPng-specific statistics for the VLAN v1:

```
show ripng interface v1
```

The following is sample output from this command:

VLAN	: v1	Interface	: 22cc::3/64
Router RIPng	: Enabled	Cost	: 15
Input Policy	: None	Output Policy	: None
Trusted GW Policy	: gw6	Poison Reverse	: Enabled
Split Horizon	: Enabled	Triggered Updates	: Enabled
Rcvd Packets	: 106358	Sent Packets	: 106358
Sent Trig. Updates	: 3	Rcvd Bad Packets	: 0
Rcvd Bad Routes	: 0		

```
Neighbor Addresses : fe80::201:30ff:fe94:f400
Interface Addresses : 22cc::3/64, fe80::280:c8ff:feb9:2855/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ripng routes

```
show ripng routes {detail} {network <ripngNetworkPrefix>}
```

Description

Displays all matching routes in the RIPng routing database.

Syntax Description

detail	Displays all available information from the RIPng routing table.
ipv6-prefix	Specifies the route prefix for the routes to show.
prefix-length	Specifies the address mask of the IPv6 prefix.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIPng, including routes exported from the system routing table and originated by other protocols, for example OSPFv3 (also called redistributed routes).

Example

The following command displays a summary of RIPng specific routes:

```
show ripng routes
```

The following is sample output from this command:

Network	Next Hop	Mtr	VLAN
*> 2aaa::/64	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	2	v2
*> 2bbb::/64	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ccc::/64	(local)	1	(direct)
*	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ddd::/64	(local)	1	(direct)
*	fe80::201:30ff:fe94:f400%v2	2	v2

The following command displays the detailed RIPng route information:

```
show ripng routes detail
```


The following is sample output from this command:

```
IPv6 RIPng routing table entry for 2aaa::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 2, tag 0, timeout in 02:44, valid

IPv6 RIPng routing table entry for 2bbb::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid

IPv6 RIPng routing table entry for 2ccc::/64
Paths: (3 available, best #1)
  Local from direct
    Metric 1, tag 0, no timeout, valid, best
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure ripng

```
unconfigure ripng {vlan <vlan-name> | tunnel <tunnel-name> | vlan all |  
tunnel all}
```

Description

Resets RIPng parameters to the default value.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies either all IPv6 configured VLANs or all IPv6 tunnels.

Default

N/A.

Usage Guidelines

Issuing the command `unconfigure ripng` resets all the interfaces and the global configuration to the defaults, and disables RIPng, as that is the default.

Example

The following command resets the RIPng configuration to the default for the VLAN *finance*:

```
unconfigure rip finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

32 OSPF Commands

This chapter describes commands used for the interior gateway protocol OSPF.

Open Shortest Path First (OSPF) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system* (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPF allows parts of a network to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other ABRs.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.



NOTE

Do not set the router ID to 0.0.0.0.

Licensing

To use the complete OSPF functionality, you must have a Core license installed on your switch. The BlackDiamond 10808 ships with a Core, or Advanced Core license. Other platforms, except the Summit X450e and X250e, can be upgraded to a Core license. A SummitStack can operate at a Core license level if Summit X450e and X250e switches are not included in the stack. See the section [“Software Licenses and Feature Packs” on page 37](#) for more information about licensing.

A subset of OSPF, called OSPF Edge Mode, is available with an Advanced Edge license.

OSPF Edge Mode

OSPF Edge Mode is a subset of OSPF available on platforms with an Advanced Edge license. There are two restrictions on OSPF Edge Mode:

- At most, four Active OSPF VLAN interfaces are permitted. There is no restriction on the number of Passive interfaces.
- The OSPF Priority on VLANs is zero, and is not configurable. This prevents the system from acting as a DR or BDR

clear ospf counters

```
clear ospf counters
{ interfaces [all | vlan <vlan-name> | area <area-identifier>]
| area [all | <area-identifier>]
| virtual-link [all | <router-identifier> <area-identifier>]
| neighbor [all | routerid [<ip-address> {<ip-mask>} | <ipNetmask>]
| vlan <vlan-name>]
| system
}
```

Description

Clears the OSPF counters (statistics).

Syntax Description

vlan-name	Specifies a VLAN name.
router-identifier	Specifies a router interface number.
area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask
system	Specifies the OSPF system counters.

Default

N/A.

Usage Guidelines

The global command `clear counters` also clears all OSPF counters. This global command is the equivalent of `clear ospf counters` for OSPF.

Example

The following command clears the OSPF counters for area 1.1.1.1:

```
clear ospf counters area 1.1.1.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf cost

```
configure ospf [area <area-identifier> | vlan [<vlan-name> | all]] cost  
[automatic | <cost>]
```

Description

Configures the cost metric of one or all interface(s) or an area.

Syntax Description

area-identifier	Specifies an OSPF area.
vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
automatic	Determine the advertised cost from the OSPF metric table.
cost	Specifies the cost metric.

Default

The default cost is automatic.

Usage Guidelines

The range is 1 through 65535.

Example

The following command configures the cost metric of the VLAN *accounting*:

```
configure ospf vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf priority

```
configure ospf [area <area-identifier> | vlan [<vlan-name> | all]] priority  
<priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPF interface(s) or for all the interfaces within the area.

Syntax Description

area-identifier	Specifies an OSPF area.
vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

The range is 0 through 255, and the default setting is 1. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospf area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a core license.

configure ospf authentication

```
configure ospf [vlan [<vlan-name> | all] | area <area-identifier> |
virtual-link <router-identifier> <area-identifier>] authentication
[{encrypted} simple-password <simple-password> | {encrypted} md5
<md5_key_id> <md5_key>| none]
```

Description

Specifies the authentication password (up to eight characters) or Message Digest 5 (MD5) key for one or all interfaces in a specific area or a virtual link.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router.
encrypted	Indicates that the password (or key) is already encrypted (do not use this option).
simple-password	Specifies an authentication password (up to 8 ASCII characters).
md5-key_id	Specifies a Message Digest 5 key, from 0-255.
md5_key	Specifies a numeric value from 0-65,536. Can also be alphanumeric, up to 26 characters.
none	Disables authentication.

Default

N/A.

Usage Guidelines

The `md5_key` is a numeric value with the range 0 to 65,536 or alphanumeric. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.

The `encrypted` option is used by the switch when generating a configuration file and when parsing a switch-generated configuration file. Do not select the `encrypted` option in the CLI.

Example

The following command configures MD5 authentication on the VLAN *subnet_26*:

```
configure ospf vlan subnet_26 authentication md5 32 test
```


History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf add virtual-link

```
configure ospf add virtual-link <router-identifier> <area-identifier>
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router interface number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0, and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link between the two interfaces:

```
configure ospf add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf add vlan area

```
configure ospf add vlan [<vlan-name> | all] area <area-identifier>
{passive}
```

Description

Enables OSPF on one or all VLANs (router interfaces).

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
area-identifier	Specifies the area to which the VLAN is assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables OSPF on a VLAN named *accounting*:

```
configure ospf add vlan accounting area 0.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf add vlan area link-type

```
configure ospf add vlan <vlan-name> area <area-identifier> link-type [auto  
| broadcast | point-to-point] {passive}
```

Description

Configures the OSPF link type.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPF link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving packets on this interface.

Default

Auto.

Usage Guidelines

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command configures the OSPF link type as automatic on a VLAN named *accounting*:

```
configure ospf add vlan accounting area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area external-filter

```
configure ospf area <area-identifier> external-filter [<policy-map> | none]
```

Description

Configures an external filter policy.

Syntax Description

area-identifier	Specifies the OSPF target area.
policy-map	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF external routes from being advertised into that area.

Using the none mode specifies that no external filter is applied.

Example

The following command configures an external filter policy, *nosales*:

```
configure ospf area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area interarea-filter

```
configure ospf area <area-identifier> interarea-filter [<policy-map> | none]
```

Description

Configures a global inter-area filter policy.

Syntax Description

area-identifier	Specifies the OSPF target area.
policy-map	Specifies a policy.
none	Specifies not to apply an interarea filter.

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF inter-area routes from being sourced from any other areas.

Example

The following command configures an inter-area filter policy, *nosales*:

```
configure ospf area 0.0.0.6 interarea-filter nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area add range

```
configure ospf area <area-identifier> add range [<ip-address> <ip-mask> |
<ipNetmask>] [advertise | noadvert] [type-3 | type-7]
```

Description

Configures a range of IP addresses in an OSPF area to be aggregated.

Syntax Description

area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvertise	Specifies not to advertise the aggregated range of IP addresses.
type-3	Specifies type 3 LSA, summary LSA.
type-7	Specifies type 7 LSA, NSSA external LSA.

Default

N/A.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address:

```
configure ospf area 1.2.3.4 add range 10.1.2.0/24 advertise type-3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area delete range

```
configure ospf area <area-identifier> delete range [<ip-address> <ip-mask>
| <ipNetmask>]
```

Description

Deletes a range of aggregated IP addresses in an OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an aggregated IP address range:

```
configure ospf area 1.2.3.4 delete range 10.1.2.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area normal

```
configure ospf area <area-identifier> normal
```

Description

Configures an OSPF area as a normal area.

Syntax Description

area-identifier	Specifies an OSPF area.

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPF area as a normal area:

```
configure ospf area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area nssa stub-default-cost

```
configure ospf area <area-identifier> nssa [summary | nosummary] stub-  
default-cost <cost> {translate}
```

Description

Configures an OSPF area as an NSSA.

Syntax Description

area-identifier	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.
translate	Specifies whether type-7 LSAs are translated into type-5 LSAs.

Default

N/A.

Usage Guidelines

NSSAs are similar to the existing OSPF stub area configuration option, but have the following two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area, if translated to type 5 LSAs.

When configuring an OSPF area as an NSSA, the translate option should only be used on NSSA border routers, where translation is to be enforced. If translate is not used on any NSSA border router in a NSSA, one of the ABRs for that NSSA is elected to perform translation (as indicated in the NSSA specification). The option should not be used on NSSA internal routers. Doing so inhibits correct operation of the election algorithm.

Example

The following command configures an OSPF area as an NSSA:

```
configure ospf area 10.1.1.0 nssa summary stub-default-cost 10 translate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area stub stub-default-cost

```
configure ospf area <area-identifier> stub [summary | nosummary] stub-  
default-cost <cost>
```

Description

Configures an OSPF area as a stub area.

Syntax Description

area-identifier	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.

Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory and computation requirements on OSPF routers.

Example

The following command configures an OSPF area as a stub area:

```
configure ospf area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf area timer

```
configure ospf area <area-identifier> timer <retransmit-interval> <transit-
delay> <hello-interval> <dead-interval> {<wait-timer-interval>}
```

Description

Configures the timers for all interfaces in the same OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.
wait-timer-interval	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- retransmit interval—Default: 5
- transit delay—Default: 1
- hello interval—Default: 10
- dead interval—Default: 40
- wait timer interval—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.
- wait timer interval—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospf area 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf ase-limit

```
configure ospf ase-limit <number> {timeout <seconds>}
```

Description

Configures the AS-external LSA limit and overflow duration associated with OSPF database overflow handling.

Syntax Description

number	Specifies the number of external routes that can be held in a link-state database.
seconds	Specifies a duration for which the system has to remain in the overflow state.

Default

The default for timeout is 0, which indicates that once the router goes into overflow state, it stays there until OSPF is disabled and then re-enabled.

Usage Guidelines

None.

Example

The following command configures the AS-external LSA limit and overflow duration:

```
configure ospf ase-limit 50000 timeout 1800
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf ase-summary add

```
configure ospf ase-summary add [<ip-address> <ip-mask> | <ipNetmask>] cost
<cost> {tag <number>}
```

Description

Aggregates AS-external routes in a specified address range.

Syntax Description

ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.
cost	Specifies a metric that will be given to the summarized route.
tag	Specifies an OSPF external route tag.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command summarizes AS-external routes:

```
configure ospf ase-summary add 175.1.0.0/16 cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf ase-summary delete

```
configure ospf ase-summary delete [<ip-address> <ip-mask> | <ipNetmask>]
```

Description

Deletes an aggregated OSPF external route.

Syntax Description

ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command deletes the aggregated AS-external route:

```
configure ospf ase-summary delete 175.1.0.0/16
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf delete virtual-link

```
configure ospf delete virtual-link <router-identifier> <area-identifier>
```

Description

Removes a virtual link.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a virtual link:

```
configure ospf delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf delete vlan

```
configure ospf delete vlan [<vlan-name> | all]
```

Description

Disables OSPF on one or all VLANs (router interfaces).

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPF on VLAN *accounting*:

```
configure ospf delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf import-policy

```
configure ospf import-policy [<policy-map> | none]
```

Description

Associates or removes the policy applied to OSPF routes added to the system routing table.

Syntax Description

policy-map	Specifies the policy to apply.

Default

No policy.

Usage Guidelines

Use this command to associate a policy with the OSPF routes installed into the system table. Use the `none` option to remove the policy association.

Example

The following example applies the policy *campuseast* to OSPF routes:

```
configure ospf import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf lsa-batch-interval

```
configure ospf lsa-batch-interval <seconds>
```

Description

Configures the OSPF LSA batching interval.

Syntax Description

seconds	Specifies a time in seconds.

Default

The default setting is 30 seconds.

Usage Guidelines

The range is between 0 (disabled) and 600 seconds, using multiples of 5 seconds. The LSAs added to the LSDB during the interval are batched together for refresh or timeout.

Example

The following command configures the OSPF LSA batch interval to a value of 100 seconds:

```
configure ospf lsa-batch-interval 100
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf metric-table

```
configure ospf metric-table 10M <cost_10m> 100M <cost_100m> 1G <cost_1g>
{10G <cost_10g>}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps interface.

Syntax Description

cost	Specifies the interface cost for the indicated interfaces.

Default

- 10 Mbps—The default cost is 10.
- 100 Mbps—The default cost is 5.
- 1 Gbps—The default cost is 4.
- 10 Gbps—The default cost is 2.

Usage Guidelines

None.

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospf metric-table 10m 20 100m 10 1g 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf restart

```
configure ospf restart [none | planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart router.

Syntax Description

none	Do not act as a graceful OSPF restart router.
planned	Only act as a graceful OSPF restart router for planned restarts.
unplanned	Only act as a graceful OSPF restart router for unplanned restarts.
both	Act as a graceful OSPF restart router for both planned and unplanned restarts.

Default

The default is none.

Usage Guidelines

This command configures the router as a graceful OSPF router. When configured for planned restarts, it will advertise Grace-LSAs before restarting (for example, during an upgrade of the OSPF module). When configured for unplanned restarts, it will advertise Grace-LSAs after restarting but before sending any Hellos. When configured for both, the router will advertise restarting regardless of whether the restart was planned or unplanned.

Example

The following command configures a router to perform graceful OSPF restarts only for planned restarts:

```
configure ospf restart planned
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure ospf restart grace-period

```
configure ospf restart grace-period <seconds>
```

Description

Configures the grace period sent out in Grace-LSAs and used by a restarting router.

Syntax Description

seconds	Grace period, in seconds. The default value is 120 seconds. Range is 1 to 1800 seconds.

Default

The default is 120 seconds.

Usage Guidelines

This command configures the grace period sent out to helper neighbor routers and used by the restarting router. The value of the grace period must be greater than the dead interval, and less than the LSA refresh time.

Example

The following command configures a router to send LSAs with a 240 second grace period during graceful OSPF restarts:

```
configure ospf restart grace-period 240
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure ospf restart-helper

```
configure ospf [vlan [all | <vlan-name>] | area <area-identifier> |  
virtual-link <router-identifier> <area-identifier>] restart-helper [none |  
planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart helper router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.
none	Do not act as a graceful OSPF restart helper router.
planned	Only act as a graceful OSPF restart helper router for planned restarts.
unplanned	Only act as a graceful OSPF restart helper router for unplanned restarts.
both	Act as a graceful OSPF restart helper router for both planned and unplanned restarts.

Default

The router default is none.

Usage Guidelines

This command configures the router as a graceful OSPF restart helper router for a single or multiple routers. When the router is acting as a helper, it will continue to advertise the restarting router as if it was fully adjacent.

One OSPF interface may not help more than one restarting router. An OSPF interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router will not support graceful restart for its neighbor.

Example

The following command configures a router to be a graceful OSPF helper router for planned restarts for all routers in area 10.20.30.40:

```
configure ospf area 10.20.30.40 restart-helper planned
```


History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure ospf routerid

```
configure ospf routerid [automatic | <router-identifier>]
```

Description

Configures the OSPF router ID. If automatic is specified, the switch uses the highest IP interface address as the OSPF router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
router-identifier	Specifies a router address.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.



NOTE

Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID:

```
configure ospf routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf spf-hold-time

```
configure ospf spf-hold-time <seconds>
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

seconds	Specifies a time in seconds. The range is 0 to 300 seconds.

Default

3 seconds.

Usage Guidelines

None.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospf spf-hold-time 6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf virtual-link timer

```
configure ospf virtual-link <router-identifier> <area-identifier> timer
<retransmit-interval> <transit-delay> <hello-interval> <dead-interval>
```

Description

Configures the timers for a virtual link.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5
- transit delay—Default: 1
- hello interval—Default: 10
- dead interval—Default: 40
- wait timer interval—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Example

The following command sets the timers on the virtual link in area 0.0.0.2 and remote router ID 6.6.6.6:

```
configure ospf virtual-link 6.6.6.6 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf vlan area

```
configure ospf vlan <vlan-name> area <area-identifier>
```

Description

Associates a VLAN (router interface) with an OSPF area. By default, all router interfaces are associated with area 0.0.0.0.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.

Default

Area 0.0.0.0

Usage Guidelines

Any OSPF network that contains more than one area is required to have an area configured as area 0, also called the *backbone*. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0, and then expand into other areas.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all networks outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

When a VLAN is configured to run OSPF, by default you must assign it to an area.

Example

The following command associates the VLAN *accounting* with an OSPF area:

```
configure ospf vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf vlan neighbor add

```
configure ospf vlan <vlan-name> neighbor add <ip-address>
```

Description

Configures the IP address of a point-to-point neighbor.

Syntax Description

vlan-name	Specifies a VLAN name.
ip-address	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor add 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf vlan neighbor delete

```
configure ospf vlan <vlan-name> neighbor delete <ip-address>
```

Description

Deletes the IP address of a point-to-point neighbor.

Syntax Description

vlan-name	Specifies a VLAN name.
ip-address	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor delete 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ospf vlan timer

```
configure ospf vlan [<vlan-name> | all] timer <retransmit-interval>
<transit-delay> <hello-interval> <dead-interval> {<wait-timer-interval>}
```

Description

Configures the OSPF wait interval for a VLAN or all VLANs.

Syntax Description

vlan-name	Specifies a VLAN name.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647.
wait-timer-interval	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- retransmit interval—5 seconds.
- transit delay—1 second.
- hello interval—10 seconds.
- dead interval—40 seconds.
- wait timer interval—dead interval.

Usage Guidelines

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.
- wait timer interval—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command configures the OSPF wait interval on the VLAN *accounting*:

```
configure ospf vlan accounting timer 10 15 20 60 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create ospf area

```
create ospf area <area-identifier>
```

Description

Creates an OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.

Default

Area 0.0.0.0

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates an OSPF area:

```
create ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete ospf area

```
delete ospf area [<area-identifier> | all]
```

Description

Deletes an OSPF area or all OSPF areas.

Syntax Description

area-identifier	Specifies an OSPF area.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPF area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPF area:

```
delete ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ospf

```
disable ospf
```

Description

Disables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the OSPF process for the router:

```
disable ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ospf capability opaque-lsa

```
disable ospf capability opaque-lsa
```

Description

Disables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command disables opaque LSAs across the entire system:

```
disable ospf capability opaque-lsa
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ospf export

```
disable ospf export [bgp | direct | e-bgp | i-bgp | rip | static]
```

Description

Disables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies direct routes.
i-bgp	Specifies I-BGP routes.
e-bgp	Specifies E-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPF from exporting routes derived from other protocols.

Example

The following command disables OSPF to export BGP-related routes to other OSPF routers:

```
disable ospf export bgp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ospf originate-default

```
disable ospf originate-default
```

Description

Disables the generation of a default external LSA.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables generating a default external LSA:

```
disable ospf originate-default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ospf restart-helper-lsa-check

```
disable ospf [vlan [all | <vlan-name>] | area <area-identifier> | virtual-link <router-identifier> <area-identifier>] restart-helper-lsa-check
```

Description

Disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.

Default

The default is enabled.

Usage Guidelines

This command disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Example

The following command disables a router from terminating graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
disable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

disable ospf use-ip-router-alert

```
disable ospf use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing OSPF control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the OSPF router alert IP option:

```
disable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ospf

```
enable ospf
```

Description

Enables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the OSPF process for the router:

```
enable ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ospf capability opaque-lsa

```
enable ospf capability opaque-lsa
```

Description

Enables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command enables opaque LSAs across the entire system:

```
enable ospf capability opaque-lsa
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ospf export

```
enable ospf export [bgp | direct | e-bgp | i-bgp | rip | static]
[cost <cost> type [ase-type-1 | ase-type-2] {tag <number>} | <policy-map>]
```

Description

Enables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
i-bgp	Specifies I-BGP routes.
direct	Specifies direct routes.
e-bgp	Specifies E-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
policy-map	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After OSPF export is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

The cost metric is inserted for all BGP, RIP-learned, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Example

The following command enables OSPF to export BGP-related routes using LSAs to other OSPF routers:

```
enable ospf export bgp cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ospf originate-default

```
enable ospf originate-default {always} cost <cost> type [ase-type-1 | ase-type-2] {tag <number>}
```

Description

Enables a default external LSA to be generated by OSPF, if no other default route is originated by OSPF by way of RIP and static route re-distribution.

Syntax Description

always	Specifies for OSPF to always advertise the default route.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.

Default

N/A.

Usage Guidelines

If always is specified, OSPF always advertises the default route. If always is not specified, OSPF adds the default LSA if a reachable default route is in the route table.

Example

The following command generates a default external type-1 LSA:

```
enable ospf originate-default cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ospf restart-helper-lsa-check

```
enable ospf [vlan [all | <vlan-name>] | area <area-identifier> | virtual-link <router-identifier> <area-identifier>] restart-helper-lsa-check
```

Description

Enables the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.

Default

The default is enabled.

Usage Guidelines

This command configures the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router. This will occur when the restart-helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Example

The following command configures a router to terminate graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
enable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

enable ospf use-ip-router-alert

```
enable ospf use-ip-router-alert
```

Description

Enables the generation of the OSPF router alert IP option.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the OSPF router alert IP option:

```
enable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf

```
show ospf
```

Description

Displays global OSPF information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPF information:

```
show ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf area

```
show ospf area {<area-identifier>}
```

Description

Displays information about OSPF areas.

Syntax Description

area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about OSPF area 1.2.3.4:

```
show ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf area detail

```
show ospf area detail
```

Description

Displays information about all OSPF areas.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF areas:

```
show ospf area detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf ase-summary

```
show ospf ase-summary
```

Description

Displays the OSPF external route aggregation configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the OSPF external route aggregation configuration:

```
show ospf ase-summary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf interfaces detail

```
show ospf interfaces detail
```

Description

Displays detailed information about all OSPF interfaces.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF interfaces:

```
show ospf interfaces detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf interfaces

```
show ospf interfaces {vlan <vlan-name> | area <area-identifier>}
```

Description

Displays information about one or all OSPF interfaces.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.

Default

If no argument is specified, all OSPF interfaces are displayed.

Usage Guidelines

None.

Example

The following command displays information about one or all OSPF interfaces on the VLAN *accounting*:

```
show ospf interfaces vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf lsdb

```
show ospf lsdb {detail | stats} {area [<area-identifier> | all]}
{{lstype} [<lstype> | all]} {lsid <lsid-address>{<lsid-mask>}}
{routerid <routerid-address> {<routerid-mask>}} {interface[[<ip-
address>{<ip-mask>} | <ipNetmask>] | vlan <vlan-name>]}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
stats	Specifies to display the number of matching LSAs, but not any of their contents.
area-identifier	Specifies an OSPF area.
all	Specifies all OSPF areas.
lstype	Specifies an LS type
lsid	Specifies an LS ID.
lsid-mask	Specifies an LS ID mask
interface	Specifies to display interface types.
routerid-address	Specifies a LSA router ID address.
vlan-name	Specifies a VLAN name.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospf lsdb` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospf lsdb
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all` with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospf lsdb
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf memory

```
show ospf memory {detail | <memoryType>}
```

Description

Displays OSPF specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPF specific memory for all types:

```
show ospf memory detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf neighbor

```
show ospf neighbor {routerid [<ip-address> {<ip-mask>} | <ipNetmask>]}  
{vlan <vlan-name>} {detail}
```

Description

Displays information about an OSPF neighbor.

Syntax Description

ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask
vlan-name	Specifies a VLAN name.
detail	Specifies detail information.

Default

If no argument is specified, all OSPF neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPF neighbors on the VLAN *accounting*:

```
show ospf neighbor vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ospf virtual-link

```
show ospf virtual-link {<router-identifier> <area-identifier>}
```

Description

Displays virtual link information about a particular router or all routers.

Syntax Description

router-identifier	Specifies a router interface number.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

area-identifier—Transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays virtual link information about a particular router:

```
show ospf virtual-link 1.2.3.4 10.1.6.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure ospf

```
unconfigure ospf {vlan <vlan-name> | area <area-identifier>}
```

Description

Resets one or all OSPF interfaces to the default settings.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets the OSPF interface to the default settings on the VLAN *accounting*:

```
unconfigure ospf accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

This chapter describes commands used for the IPv6 interior gateway protocol OSPFv3.

Open Shortest Path First (OSPFv3) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system* (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router in an area has an identical database maintained from the perspective of that router.

OSPFv3 supports IPv6, and uses commands only slightly modified from that used to support IPv4. OSPFv3 has retained the use of the four-byte, dotted decimal numbers for router IDs, LSA IDs, and area IDs.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPFv3 allows parts of a network to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPFv3 are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other OSPFv3 routers.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPFv3 and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPFv3 must have a unique router ID. It is recommended especially that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address, since your router may not have an IPv4 address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.

**NOTE**

Do not set the router ID to 0.0.0.0.

Licensing

To use OSPFv3, you must have a Core license installed on your switch. The BlackDiamond 10808 ships with a Core, or Advanced Core license. Other platforms can be upgraded to a Core license. See the section [“Software Licenses and Feature Packs” on page 37](#) for more information about licensing.

clear ospfv3 counters

```
clear ospfv3 {domain <domainName>} counters
{ interfaces [[vlan | tunnel] all | vlan <vlan-name> |
  tunnel <tunnel-name> | area <area-identifier>]
| area [all | <area-identifier>]
| virtual-link [all | {routerid} <router-identifier>
  {area} <area-identifier>]
| neighbor [all | routerid <router-identifier> | vlan <vlan-name> |
  tunnel <tunnel-name>]
| system
}
```

Description

Clears the OSPFv3 counters (statistics).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all VLANs, tunnels, areas, neighbors, or virtual-links.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
system	Specifies the OSPFv3 system/global counters.

Default

N/A.

Usage Guidelines

The global command `clear counters` also clears all OSPFv3 counters. This global command is the equivalent of `clear ospfv3 counters` for OSPFv3.

This command can be used to clear various OSPFv3 counters (Interface, Area, Virtual-Link, System etc.). The following is the list of various counters that would be reset to zero by this command:

- Neighbor specific counters
 - Number of state changes
 - Number of events
- Interface/VLAN/Virtual-link/Tunnel specific counters
 - Number of Hellos rxed
 - Number of Hellos txed
 - Number of DB Description rxed

- Number of DB description txed
- Number of LS request rxed
- Number of LS request txed
- Number of LS update rxed
- Number of LS update txed
- Number of LS ack rxed
- Number of LS ack txed
- Number of rxed OSPFv3 packet discarded
- Number of state changes
- Number of events
- Area Specific counters
 - All counters of interfaces associated with an area
 - Number of SPF runs
- Domain (global)/system specific counters
 - Number of self originated LSAs
 - Number of received LSAs

Example

The following command clears the OSPFv3 counters for area 1.1.1.1:

```
clear ospfv3 counters area 1.1.1.1
```

The following command clears all the OSPFv3 counters for the neighbor 192.168.0.1 in the domain *ospf-core*:

```
clear ospfv3 domain ospf-core counters neighbor routerid 192.168.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 add interface

```
configure ospfv3 {domain <domainName>} add [vlan <vlan-name> | tunnel
<tunnel-name>] {instance-id <instanceId>} area <area-identifier> link-type
[auto | broadcast | point-to-point] {passive}
```

Description

Enables OSPFv3 on an interface.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
instanceId	Specifies the instance ID for this interfaces. Range is 0 to 255.
area-identifier	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPFv3 link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

The default link-type is Auto.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on an IPv6 configured VLAN or an IPv6 tunnel. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Enable IPv6 forwarding before enabling OSPFv3, otherwise, you will receive a warning message.

**NOTE**

Configuration of the link-type parameter is not supported. OSPFv3 will always consider the link-type to be broadcast.

Example

The following command adds the VLAN *accounting* (enabling OSPFv3 on the interface), to the area 0.0.0.1 with an instance ID of 2:

```
configure ospfv3 add vlan accounting instance-id 2 area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 add interface all

```
configure ospfv3 {domain <domainName>} add [vlan | tunnel] all {instance-id
<instanceId>} area <area-identifier> {passive}
```

Description

Enables OSPFv3 on all VLANs or all tunnels (router interfaces).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
instanceId	Specifies the instance ID for these interfaces. Range is 0 to 255.
area-identifier	Specifies the area to which the interfaces are assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

OSPFv3 is disabled on the interfaces.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on all IPv6 configured VLANs or all IPv6 tunnels. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command enables OSPFv3 on all IPv6 tunnels:

```
configure ospfv3 add tunnel all area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 add virtual-link

```
configure ospfv3 {domain <domainName>} add virtual-link {routerid} <router-identifier> {area} <area-identifier>
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router identifier, a four-byte, dotted decimal number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area add range

```
configure ospfv3 {domain <domainName>} area <area-identifier> add range
<ipv6netmask> [advertise | noadvert] inter-prefix
```

Description

Configures a range of IP addresses in an OSPFv3 area to be aggregated.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
ipv6netmask	Specifies an IPv6 address / prefix length.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvert	Specifies not to advertise the aggregated range of IP addresses.
inter-prefix	Specifies aggregate, inter-area-prefix LSAs.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address to area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 add range 2aaa:456:3ffe::/64 advertise inter-prefix
```

History

This command was first available in ExtremeXOS 11.2

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area cost

```
configure ospfv3 {domain <domainName>} area <area-identifier> cost  
[automatic | <cost>]
```

Description

Configures the cost of sending a packet to all interfaces belonging to an area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
automatic	Determine the advertised cost from the OSPFv3 metric table.
cost	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic. The default domain is OSPF-Default.

Usage Guidelines

Use this command to set the cost of the links belonging to area manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost of area 0.0.0.1 to 10. All the links of this area will inherit the area's cost value of 10.

```
configure ospfv3 domain ospf-enterprise area 0.0.0.1 cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area delete range

```
configure ospfv3 {domain <domainName>} area <area-identifier> delete range
<ipv6netmask>
```

Description

Removes a range of IP addresses in an OSPFv3 area to be aggregated.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
ipv6netmask	Specifies an IPv6 address / prefix length.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If you attempt to delete a range that was not configured, you will receive an error message.

Example

The following command is used to delete a summary network from area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 delete range 2aaa:456:3ffe::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area external-filter

```
configure ospfv3 {domain <domainName>} area <area-identifier> external-
filter [<policy-map> | none]
```

Description

Configures an external filter policy.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies the OSPFv3 target area.
policy-map	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPFv3 areas (an ABR function), a policy can be applied to an OSPFv3 area that filters a set of OSPFv3 external routes from being advertised into that area, in other words, filtering some of the inbound AS-external-LSAs.

OSPFv3 routers that do not have enough memory to hold the entire AS-external-LSAs should configure an external area filter to drop part of the external-LSAs. Configuring this policy will enable routers with limited resources to be put into an OSPFv3 network.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
- Action (set) attributes
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external area policy file:

```
entry one {
    if match any{
        nlri 2001:db8:3e5c::/48;
        nlri 2001:db8:2146:2341::/64;
    } then {
        deny;
    }
}
```

Example

The following command configures an external filter policy, *nosales* for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area interarea-filter

```
configure ospfv3 {domain <domainName>} area <area-identifier> interarea-
filter [<policy-map> | none]
```

Description

Configures an inter-area filter policy.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies the OSPFv3 target area.
policy-map	Specifies a policy.
none	Specifies not to apply an inter-area filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 can apply an inter-area policy to filter some inter-area-prefix-LSAs and inter-area-router-LSAs from other areas. This can reduce the size of link state database of routers belonging to the area.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
- Action (set) attributes
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external area policy file:

```
entry one {
  if match any{
    nlri 2001:db8:3e5c::/48;
    nlri 2001:db8:2146:2341::/64;
  } then {
    deny;
  }
}
```

```
entry two {  
    if match any{  
        nlri 2001:db8:444::/48;  
        nlri 2001:db8:541f:65bd::/64;  
    } then {  
        permit;  
    }  
}
```

Example

The following command configures an inter-area filter policy, *nosales* for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 interarea-filter nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area normal

```
configure ospfv3 {domain <domainName>} area <area-identifier> normal
```

Description

Configures an OSPFv3 area as a normal area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPFv3 area as a normal area:

```
configure ospfv3 area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area priority

```
configure ospfv3 {domain <domainName>} area <area-identifier> priority  
<priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for all the interfaces within the area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospfv3 area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area stub

```
configure ospfv3 {domain <domainName>} area <area-identifier> stub [summary  
| nosummary] stub-default-cost <cost>
```

Description

Configures an OSPFv3 area as a stub area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
summary	Specifies that inter-area LSAs can be propagated into the area.
nosummary	Specifies that inter-area LSAs cannot be propagated into the area.
cost	Specifies a cost metric.

Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory consumption requirements on OSPFv3 routers.

Example

The following command configures an OSPFv3 area as a stub area:

```
configure ospfv3 area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 area timer

```
configure ospfv3 {domain <domainName>} area <area-identifier> timer
{retransmit-interval} <retransmit-interval> {transit-delay} <transit-delay>
{hello-interval} <hello-interval> {dead-interval} <dead-interval>
```

Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

Configuring OSPFv3 timers on a per-area basis is a shorthand for applying the timers to each VLAN and tunnel in the area at the time of configuration. If you add more VLANs or tunnels to the area, you must configure the timers for them explicitly.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello

packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.

**NOTE**

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospfv3 area 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 delete interface

```
configure ospfv3 {domain <domainName>} delete [vlan <vlan-name> | tunnel  
<tunnel-name> | [vlan | tunnel] all]
```

Description

Disables OSPFv3 on one or all VLANs or tunnels (router interfaces).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all VLANs, or tunnels.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 on VLAN *accounting*:

```
configure ospfv3 delete vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 delete virtual-link

```
configure ospfv3 {domain <domainName>} delete virtual-link {routerid}
<router-identifier> {area} <area-identifier>
```

Description

Deletes a virtual link connected to another ABR.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router identifier, a four-byte, dotted decimal number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command deletes a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 import-policy

```
configure ospfv3 {domain <domainName>} import-policy [<policy-map> | none]
```

Description

Associates or removes the policy applied to OSPFv3 routes added to the system routing table.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
policy-map	Specifies the policy to apply.
none	Specifies to remove any policy previously applied.

Default

No policy.

Usage Guidelines

Use this command to associate a policy with the OSPFv3 routes installed into the system table. Use the `none` option to remove the policy association. The import policy cannot be used to filter routes from being added to the routing table. It can only be used to change the attributes of routes.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
 - route-origin [ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra]
- Action (set) attributes
 - cost <cost>
 - tag <number>

Any other policy attribute will not be recognized and will be ignored.

Example

The following example applies the policy *campuseast* to OSPFv3 routes:

```
configure ospfv3 import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 interface area

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name>] area <area-identifier>
```

Description

Moves an interface from one OSPFv3 area to another.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0

Usage Guidelines

Use this command to move an already configured interface from one area to another. The instance ID associated with the interface will be unchanged.

Example

The following command moves the VLAN *accounting* to the OSPFv3 area 0.0.0.6:

```
configure ospfv3 vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 interface cost

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name> | [vlan | tunnel] all]} cost [automatic | <cost>]
```

Description

Configures the cost of one or all interface(s).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
automatic	Determine the advertised cost from the OSPFv3 metric table.
cost	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic.

Usage Guidelines

Use this command to set the cost of an interface (a VLAN or tunnel) manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost metric of the VLAN *accounting*:

```
configure ospfv3 vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 interface priority

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-  
name> | [vlan | tunnel] all] priority <priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPFv3 interface(s).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets the priority of the interface VLAN corporate to 10:

```
configure ospfv3 domain ospf-internal vlan corporate priority 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 interface timer

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-
name> | [vlan | tunnel] all] timer {retransmit-interval} <retransmit-
interval> {transit-delay} <transit-delay> {hello-interval} <hello-interval>
{dead-interval} <dead-interval>
```

Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

Use this command to configure the OSPFv3 timers on a per-interface basis.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.

**NOTE**

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers for the VLAN *corporate*:

```
configure ospfv3 domain ospf-default vlan corporate timer retransmit-interval 10  
transit-delay 2 hello-interval 20 dead-interval 80
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 metric-table

```
configure ospfv3 {domain <domainName>} metric-table 10M <cost_10m> 100M
<cost_100m> 1G <cost_1g> {10G <cost_10g>}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps interface.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
cost_x	Specifies the interface cost for the indicated interfaces. Range is 1 to 65535.

Default

- 10 Mbps—The default cost is 100.
- 100 Mbps—The default cost is 50.
- 1 Gbps—The default cost is 40.
- 10 Gbps—The default cost is 20.

Usage Guidelines

The value of the costs cannot be greater for higher speed interfaces. In other words, the following condition must be true:

cost_10m >= cost_100m >= cost_1g >= cost_10g

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospfv3 metric-table 10m 200 100m 100 1g 20
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 routerid

```
configure ospfv3 {domain <domainName>} routerid [automatic | <router-
identifier>]
```

Description

Configures the OSPFv3 router ID. If automatic is specified, the switch uses the highest IPv4 interface address as the OSPFv3 router ID.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
automatic	Specifies to use automatic addressing.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPFv3 must have a unique router ID. The router ID is a four-byte, dotted decimal number, like an IPv4 address. Even though the IP address format has changed from IPv4 to IPv6, the router ID format has not. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address (if it exists). Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.

This command is accepted only when OSPFv3 is globally disabled.



NOTE

Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID to 10.1.6.1:

```
configure ospfv3 routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 spf-hold-time

```
configure ospfv3 {domain <domainName>} spf-hold-time <seconds>
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
seconds	Specifies a time in seconds. The range is 0 to 300 seconds.

Default

3 seconds.

Usage Guidelines

Setting the interval too high will force OSPFv3 to run SPF calculations less frequently. This will reduce the CPU load, but will cause delay in routes getting updated in the IP routing table. Setting the interval too low will decrease the interval between SPF calculations, but will increase the processing load on CPU.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospfv3 spf-hold-time 6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure ospfv3 virtual-link timer

```
configure ospfv3 {domain <domainName>} virtual-link {routerid} <router-
identifier> {area} <area-identifier> timer {retransmit-interval}
<retransmit-interval> {transit-delay} <transit-delay> {hello-interval}
<hello-interval> {dead-interval} <dead-interval>
```

Description

Configures the timers for a virtual link.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

In OSPFv3, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

The smaller the hello interval, the faster topological changes will be detected, but more routing traffic will ensue.

The setting of the retransmit interval should be conservative, or needless retransmissions will result. The value should be larger for serial lines and virtual links.

The transmit delay value should take into account the transmission and propagation delays for the interface.



NOTE

The wait interval is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers on the virtual link to router 6.6.6.6 transiting area 0.0.0.2:

```
configure ospfv3 virtual-link 6.6.6.6 area 0.0.0.2 timer 10 transit-delay 1  
hello-interval 20 dead-interval 200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

create ospfv3 area

```
create ospfv3 {domain <domainName>} area <area-identifier>
```

Description

Creates an OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates a non-backbone OSPFv3 area:

```
create ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

delete ospfv3 area

```
delete ospfv3 {domain <domainName>} area [<area-identifier> | all]
```

Description

Deletes an OSPFv3 area or all OSPFv3 areas.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPFv3 area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPFv3 area:

```
delete ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable ospfv3

```
disable ospfv3 {domain <domainName>}
```

Description

Disables OSPFv3 for the router.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 for the router:

```
disable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable ospfv3 export

```
disable ospfv3 {domain <domainName>} export [direct | ripng | static]
```

Description

Disables redistribution of routes to OSPFv3.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIP routes.
static	Specifies static routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPFv3 from exporting routes derived from other protocols.

Example

The following command disables OSPFv3 to export RIPng routes to other OSPFv3 routers:

```
disable ospfv3 export ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable ospfv3

```
enable ospfv3 {domain <domainName>}
```

Description

Enables OSPFv3 for the router.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.

Default

N/A.

Usage Guidelines

When OSPFv3 is enabled, it will start exchanging Hellos on all of its active interfaces. It will also start exporting routes into OSPFv3 routing domain from other protocols, if enabled.

When OSPFv3 is disabled, it will release all the run-time allocated resources like adjacencies, link state advertisements, run-time memory, etc.

OSPFv3 can be enabled successfully if and only if:

- At least one of the VLANs in the current virtual router has one IPv4 address configured

—OR—

- You explicitly configure the OSPFv3 router ID, a four-byte, dotted decimal number

Example

The following command enables OSPFv3 for the router:

```
enable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable ospfv3 export

```
enable ospfv3 {domain <domainName>} export [direct | ripng | static]
[cost <cost> type [ase-type-1 | ase-type-2] | <policy-map>]
```

Description

Enables redistribution of routes to OSPFv3.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIPng routes.
static	Specifies static routes.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
policy-map	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into OSPFv3. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.



NOTE

Setting the tag value is not supported in this release.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>

- Action (set) attributes
 - cost <cost>
 - tag <number>
 - cost-type [ase-type-1 | ase-type-2]
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example OSPFv3 export policy file:

```
entry first {
  if match any{
    nlri 2001:db8:200:300::/64;
    nlri 2001:db8:2146:23d1::/64;
    nlri 2001:db8:af31:3d0::/64;
    nlri 2001:db8:f6:2341::/64;
  } then {
    deny;
  }
}
entry second {
  if match any{
    nlri 2001:db8:304::/48;
    nlri 2001:db8:ca11::/48;
    nlri 2001:db8:da36::/48;
    nlri 2001:db8:f6a6::/48;
  } then {
    cost 220;
    cost-type ase-type-2;
    permit;
  }
}
```

Example

The following command enables OSPFv3 to export RIPng-related routes and associates a policy *redist*:

```
enable ospfv3 export ripng redist
```

History

This command was first available in ExtremeXOS 11.2.

The `tag` keyword was removed in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3

```
show ospfv3 {domain <domainName>}
```

Description

Displays global OSPFv3 information.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPFv3 information:

```
show ospfv3
```

The following is sample output:

```

OSPF Domain Name      : OSPF-Default
OSPF                  : Enabled           RouterId          : 20.0.0.1
RouterId Selection    : Automatic         ASBR                : Yes
ABR                   : Yes               ExtLSAs             : 3
ExtLSAChecksum        : 0x19420          OriginateNewLSAs    : 37
ReceivedNewLSAs       : 12              SpfHoldTime         : 10s
Num of Areas          : 5                10M Cost            : 100
100M Cost             : 50               1000M Cost (1G)     : 40
10000M Cost (10G)     : 20              Router Alert        : Disabled
ASExternal LSALimit   : Disabled         Timeout (Count)     : Disabled (0)
Originate Default     : Disabled
Import Policy File    : none

Redistribute:
  Protocol  Status   Cost   Type  Tag   Policy
  direct    Disabled 20     2     0     none
  ripng     Disabled 20     2     0     none
  static    Enabled  20     2     0     IPv6

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 area

```
show ospfv3 {domain <domainName>} area {<area-identifier> | detail}
```

Description

Displays information about OSPFv3 areas.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary information about the OSPFv3 areas:

```
show ospfv3 area
```

The following is sample output:

AREA ID	Type	Summ	Def	Num	Num	Num	SPF	Num	LSA
			Metric	ABR	ASBR	Intf	Runs	LSAs	Checksum
0.0.0.0	NORM	----	-----	0	0	1	7	7	0x3155b
1.0.0.0	NORM	----	-----	1	1	1	6	9	0x4793d
2.0.0.0	NORM	----	-----	0	0	1	5	10	0x47174
3.0.0.0	NORM	----	-----	1	0	1	3	12	0x420cf
5.0.0.0	NORM	----	-----	1	0	1	4	10	0x3b5b1

The following command displays information about OSPFv3 area 1.0.0.0:

```
show ospfv3 area 1.0.0.0
```

The following is sample output:

Area Identifier	: 1.0.0.0	Type	: NORM
Router ID	: 20.0.0.1	Num of Interfaces	: 1
Spf Runs	: 6	Num ABRs	: 1
Num ASBRs	: 1	Num DC-Bit LSAs	: 1
Num Indication LSAs	: 1	Num of DoNotAge LSAs	: 1

```

Num LSAs           : 9                LSA Chksum           : 0x4793d
Num of Nbrs        : 1                Num of Virtual Nbrs : 0

Interfaces:
Interface Name      Ospf State   DR ID           BDR ID
to65                E      BDR          0.0.0.65        20.0.0.1
accounts            E      DR           80.0.0.5         0.0.0.0
finance             E      BDR          90.0.0.7         66.0.0.4
engineering          E      ODR          192.168.0.1      165.0.0.3
Corporate           E      ODR          201.0.16.6       204.0.0.1
Inter-Area route Filter: ospfSummPolicy
External route Filter: ospfExtPolicy
Configured Address Ranges:
  Addr: fffe:408:1449::/48 Type: 3 Advt: Yes
  Addr: ffe0:930:2781::/40 Type: 7 Advt: No

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 interfaces

```
show ospfv3 {domain <domainName>} interfaces {vlan <vlan-name> | tunnel
<tunnel-name> | area <area-identifier> | detail}
```

Description

Displays information about one or all OSPFv3 interfaces.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

If no argument is specified, all OSPFv3 interfaces are displayed.

Usage Guidelines

None.

Example

The following command shows a summary of the OSPFv3 interfaces:

```
show ospfv3 interfaces
```

The following is sample output from the command:

VLAN/Tunnel	IPv6 Address	AREA ID	Flags	Cost	State	
Neighbors						
ixia	4:5:6:7::1/64	2.0.0.0	-rif	5/A	DR	1
to-jmpr	111:222:333::7/48	5.0.0.0	-rif	5/A	BDR	1
to-Zebra	3ffe:506::4/48	0.0.0.0	-rif	5/A	BDR	1
to5	234:567::7/48	3.0.0.0	-rif	5/A	BDR	1
to65	10:203:134:7::7/48	1.0.0.0	-rif	5/A	BDR	1

Flags : (f) Interface Forwarding Enabled, (i) Interface OSPF Enabled,
(p) Passive Interface, (r) Router OSPF Enable.

The following command displays information about the OSPFv3 interfaces on the VLAN *to5*:

```
show ospfv3 interfaces vlan to5
```

The following is sample output:

```
Interface           : to5                Enabled           : ENABLED
Router              : ENABLED             AreaID           : 3.0.0.0
RouterID            : 20.0.0.1            Link Type        : broadcast
Passive             : No                  Cost             : 40A
Priority             : 1                   Transit Delay     : 1s
Hello Interval      : 10s                 Rtr Dead Time    : 40s
Retransmit Interval : 5s                 Wait Timer       : 40s
Interface ID        : 63                  Instance ID      : 0
State               : BDR                 Number of state chg : 2
Hello due in        : 3s                 Number of events  : 3
Total Num of Nbrs   : 1                  Nbrs in FULL State : 1
Hellos Rxed         : 94                 Hellos Txed      : 94
DB Description Rxed  : 4                  DB Description Txed : 3
LSA Request Rxed    : 1                  LSA Request Txed  : 1
LSA Update Rxed     : 8                  LSA Update Txed   : 7
LSA Ack Rxed        : 6                  LSA Ack Txed      : 5
In Discards         : 0
DR RtId             : 10.0.0.5            BDR RtId         : 20.0.0.1
DR Interface addr    : fe80::280:c8ff:feb9:1cf1
BDR Interface addr   : fe80::280:c8ff:feb9:2089
```

Neighbors:

```
RtrId: 10.0.0.5  IpAddr: fe80::280:c8ff:feb9:1cf1  Pri: 1  Type: Auto
State: FULL  DR: 10.0.0.5  BDR: 20.0.0.1  Dead Time: 00:00:36
Options: 0x13 (-|R|-|-|E|V6)  Opaque LSA: No
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 lsdb

```
show ospfv3 {domain <domainName>} lsdb {detail} {area [<area-identifier> |
all] {lstype [router | network | inter-prefix | inter-router | intra-
prefix]} | [vlan [<vlan-name> | all] | tunnel [<tunnel-name> | all]]
{lstype link} | lstype [as-external | router | network | inter-prefix |
inter-router | intra-prefix | link]} {lsid <lsid-address>} {adv-router
<router-identifier>}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
detail	Specifies to display all fields of matching LSAs in a multi-line format.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
as-external	AS External LSA
lsid-address	Specifies the link state ID of the LSA.
routerid-identifier	Specifies the router identifier of the advertising router.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsdb
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all` with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 lsdb stats

```
show ospfv3 {domain <domainName>} lsdb stats {area [<area-identifier> |
all] {lstype [router | network | inter-prefix | inter-router | intra-
prefix]} | [vlan [<vlan-name> | all] | tunnel [<tunnel-name> | all]}
{lstype link} | lstype [as-external | router | network | inter-prefix |
inter-router | intra-prefix | link]} {lsid <lsid-address>} {adv-router
<router-identifier>}
```

Description

Displays a table of the current Link-State Database (LSDB) statistics.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
as-external	AS External LSA
lsid-address	Specifies the link state ID of the LSA.
routerid-identifier	Specifies the router identifier of the advertising router.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb stats` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsdb stats
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all`.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb stats
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 memory

```
show ospfv3 memory {detail | <memoryType>}
```

Description

Displays OSPFv3 specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPFv3 specific memory for all types:

```
show ospfv3 memory detail
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 neighbor

```
show ospfv3 {domain <domainName>} neighbor {routerid <ip-address>} {vlan  
<vlan-name> | tunnel <tunnel-name>} {detail}
```

Description

Displays information about an OSPFv3 neighbor.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
ip-address	Specifies a neighbor router ID.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
detail	Specifies detail information.

Default

If no argument is specified, all OSPFv3 neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPFv3 neighbors on the VLAN *accounting*:

```
show ospfv3 neighbor vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show ospfv3 virtual-link

```
show ospfv3 {domain <domainName>} virtual-link {{routerid} <router-  
identifier> {area} <area-identifier>}
```

Description

Displays virtual link(s) information.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

router-identifier—Router ID for the other end of the link.

area-identifier—Transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays information about the virtual link to a particular router:

```
show ospfv3 virtual-link 1.2.3.4 10.1.6.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

unconfigure ospfv3

```
unconfigure ospfv3 {domain <domainName>} {area <area-identifier> | vlan
<vlan-name> | tunnel <tunnel-name>}
```

Description

Resets one or all OSPFv3 interfaces to the default settings.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 allows you to change certain configurable OSPFv3 parameters on the fly. This command selectively resets the configurable parameters to their default values. The following is the list of parameters whose values will be reset to their defaults:

- Interface
 - Hello Interval
 - Dead Interval
 - Transmit Delay
 - Retransmit Interval
 - Priority
 - Cost
- Area
 - All the parameters of Interfaces associated with this area
 - Inter-Area-Prefix-LSA Filter
 - AS-External-LSA Filter
- OSPF Global
 - All parameters of all areas in this OSPF domain
 - SPF Delay interval
 - Interface Cost metric Table
 - Route Redistribution

Example

The following command resets the OSPFv3 interface to the default settings on the VLAN *accounting*:

```
unconfigure ospfv3 accounting
```

The following command unconfigures the parameters of the area 0.0.0.1 (and all its associated interfaces):

```
unconfigure ospfv3 domain ospf-default area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

34 BGP Commands

Border Gateway Protocol (BGP) is an exterior routing protocol that was developed for use in TCP/IP networks. The primary function of BGP is to allow different autonomous systems (ASs) to exchange network reachability information.

An autonomous system is a set of routers that are under a single technical administration. This set of routers uses a different routing protocol (such as OSPF) for intra-AS routing. One or more routers in the AS are configured to be border routers, exchanging information with other border routers (in different autonomous systems) on behalf of all of the intra-AS routers.

BGP can be used as an exterior border gateway protocol (EBGP), or it can be used within an AS as an interior border gateway protocol (IBGP).

Licensing

BGP requires a Core license, at a minimum. The BGP process will not spawn without the required license level. The MSM-1XL is shipped with an Advanced Core license and the MSM-1 is shipped with a Core license. Other platforms can be upgraded to a Core license. See the section [“Software Licenses and Feature Packs”](#) in chapter 1, [“ExtremeXOS Overview”](#), in the *ExtremeXOS Concepts Guide* for more information about licensing.

BGP Attributes

The following BGP attributes are supported by the switch:

- **Origin**—Defines the origin of the route. Possible values are IGP, EGP, and incomplete.
- **AS_Path**—The list of ASs that are traversed for this route.
- **Next_hop**—The IP address of the next hop BGP router to reach the destination listed in the NLRI field.
- **Multi_Exist_Discriminator**—Used to select a particular border router in another AS when multiple border routers exist.
- **Local_Preference**—Used to advertise this router’s degree of preference to other routers within the AS.
- **Atomic_aggregate**—Indicates that the sending border router is used a route aggregate prefix in the route update.
- **Aggregator**—Identifies the BGP router AS number and IP address that performed route aggregation.
- **Community**—Identifies a group of destinations that share one or more common attributes.
- **Cluster_ID**—Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster. A route can contain a sequence of CLUSTER_ID values representing the reflection path that the route has passed.
- **Originator_ID**—Specifies the Router_ID of the originator of the route in the local AS.

BGP Communities

A BGP community is a group of BGP destinations that require common handling. ExtremeXOS supports the following well-known BGP community attributes:

- no-export
- no-advertise
- no-export-subconfed

BGP Features

This section lists BGP features supported by ExtremeXOS:

- Route Reflectors
- Route Confederations
- Route Aggregation
- Using the Loopback Interface
- BGP Peer Groups
- BGP Route Flap Dampening
- Route Redistribution
- Policy Filtering
- Maximum Prefix Limit
- TCP MD5 Authentication
- EBGp Multihop
- Removal of private AS-Number from AS-path of outbound BGP routes
- Neighbor/Peer Group soft-reconfiguration
- BGP Graceful Restart

**NOTE**

Although the CLI commands are available, this release of ExtremeXOS does not support the MBGP/Route-refresh features.

clear bgp neighbor counters

```
clear bgp neighbor [<remoteaddr> | all] counters
```

Description

Resets the BGP counters for one or all BGP neighbor sessions to zero.

Syntax Description

remoteaddr	Specifies the IP address of a specific BGP neighbor.
all	Specifies that counters for all BGP neighbors should be reset.

Default

N/A.

Usage Guidelines

This command resets the following counters:

- In-total-msgs
- Out-total-msgs
- In-updates
- Out-updates
- Last-error
- FsmTransitions

The command `clear counters` will also reset all counter for all BGP neighbors. For BGP, the `clear counters` command is equivalent to the following BGP command:

```
clear bgp neighbor all counters
```

Example

The following command resets the counters for the BGP neighbor at 10.20.30.55:

```
clear bgp neighbor 10.20.30.55 counters
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

clear bgp flap-statistics

```
clear bgp {neighbor} <remoteaddr> {address-family [ipv4-unicast | ipv4-
multicast]} flap-statistics [all | as-path <path expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}]
```

Description

Clears flap statistics for routes to specified neighbors.

Syntax Description

all	Specifies flap statistics for all routes.
remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies a community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to clear flap statistics for a specified BGP neighbor.

The option `network any / <netMaskLen>` will clear the statistics for all BGP routes whose mask length is equal to or greater than <maskLength>, irrespective of their network address.

The option `network any / <netMaskLen> exact` will clear the statistics for all BGP routes whose mask length is exactly equal to <maskLength>, irrespective of their network address.

Example

The following command clears the flap statistics for a specified neighbor:

```
clear bgp neighbor 10.10.10.10 flap-statistics all
```

History

This command was first available in ExtremeXOS 10.1.

The `any` / `<netMaskLen>` options were added in ExtremeXOS 11.0

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp add aggregate-address

```
configure bgp add aggregate-address {address-family [ipv4-unicast |
ipv4-multicast]} <ipaddress> {as-match | as-set} {summary-only}
{advertise-policy <policy>} {attribute-policy <policy>}
```

Description

Configures a BGP aggregate route.

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
ipaddress	Specifies an IP network address and mask length.
as-match	Generates autonomous system sequence path information (order of AS numbers in AS_PATH is preserved).
as-set	Generates autonomous system set path information (order of AS numbers in AS_PATH is not preserved).
summary-only	Specifies to send only aggregated routes to the neighbors.
advertise-policy	Specifies the policy used to select routes for this aggregated route.
attribute-policy	Specifies the policy used to set the attributes of the aggregated route.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:

- 1 Enable aggregation using the following command:

```
enable bgp aggregation
```

- 2 Create an aggregate route using the following commands:

```
configure bgp add aggregate-address {address-family [ipv4-unicast | ipv4-
multicast]} <ipaddress> {as-match | as-set} {summary-only} {advertise-policy
<policy>} {attribute-policy <policy>}
```

Example

The following command configures a BGP aggregate route:

```
configure bgp add aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp add confederation-peer sub-AS-number

```
configure bgp add confederation-peer sub-AS-number <number>
```

Description

Adds a sub-AS to a confederation.

Syntax Description

number	Specifies a sub-AS number of the confederation. The range is 1 to 65535.

Default

N/A.

Usage Guidelines

Invoke this command multiple times to add multiple sub-ASs.

IBGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, all BGP speakers in each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Example

The following command adds one sub-AS to a confederation:

```
configure bgp add confederation-peer sub-AS-number 65002
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp add network

```
configure bgp add network {address-family [ipv4-unicast | ipv4-multicast]}
<ipaddr>/<mask_len> {network-policy <policy>}
```

Description

Adds a network to be originated from this router.

Syntax Description

address-family	The address family to which the network routes will be exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
ipaddr	Specifies an IP network address.
mask_len	Specifies a netmask length.
policy-name	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The network must be present in the routing table.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command adds a network to be originated from this router:

```
configure bgp add network 192.1.1.16/32
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp AS-number

```
configure bgp AS-number <number>
```

Description

Changes the local AS number used by BGP.

Syntax Description

number	Specifies a local AS number. The range is 1 to 65535.

Default

N/A.

Usage Guidelines

BGP must be disabled before the AS number can be changed.

Example

The following command changes the local AS number used by BGP:

```
configure bgp AS-number 65001
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp cluster-id

```
configure bgp cluster-id <cluster-id>
```

Description

Configures the local cluster ID.

Syntax Description

cluster-id	Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster. The range is 0 - 4294967295.

Default

N/A.

Usage Guidelines

Used when multiple route reflectors are used within the same cluster of clients.

Extreme Networks recommends disabling BGP before configuring the cluster ID.

Example

The following command appends a BGP route reflector cluster ID to the cluster list of a route:

```
configure bgp cluster-id 40000
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp confederation-id

```
configure bgp confederation-id <number>
```

Description

Specifies a BGP routing confederation ID.

Syntax Description

confederation-id	Specifies a routing confederation identifier. The range is 1 to 65535.

Default

N/A.

Usage Guidelines

IBGP requires that networks use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Use a confederation ID of 0 to indicate no confederation.

Example

The following command specifies the BGP routing confederation ID as 200:

```
configure bgp confederation-id 200
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp delete aggregate-address

```
configure bgp delete aggregate-address {address-family [ipv4-unicast |
ipv4-multicast]} [<ip address/masklength> | all]
```

Description

Deletes one or all BGP aggregated route.

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
ip address/mask length	Specifies an IP network address and netmask length.
all	Specifies all aggregated routes.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Example

The following command deletes a BGP aggregate route:

```
configure bgp delete aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp delete confederation-peer sub-AS-number

```
configure bgp delete confederation-peer sub-AS-number <number>
```

Description

Specifies a sub-AS that should be deleted from a confederation.

Syntax Description

sub-AS-number	Specifies a sub-AS.

Default

N/A.

Usage Guidelines

BGP requires that networks use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Example

The following command deletes a sub-AS from a confederation:

```
configure bgp delete confederation-peer sub-AS-number 65002
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp delete network

```
configure bgp delete network {address-family [ipv4-unicast | ipv4-
multicast]} [all | <ipaddress/mask length>]
```

Description

Deletes a network to be originated from this router.

Syntax Description

address-family	The address family to which the network routes will be exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
all	Specifies all networks.
ipaddress	Specifies an IP network address and a netmask length.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

None.

Example

The following command deletes a network to be originated from this router:

```
configure bgp delete network 192.1.1.12/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp export shutdown-priority

```
configure bgp export [direct | ospf | ospf-extern1 | ospf-extern2 | ospf-
inter | ospf-intra | rip | static] {address-family [ipv4-unicast | ipv4-
multicast]} shutdown-priority <number>
```

Description

Configures the shutdown priority for IGP export.

Syntax Description

direct	Specifies direct routing.
ospf	Specifies OSPF routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-intra	Specifies OSPF-intra routing.
rip	Specifies RIP routing.
static	Specifies static routing.
address-family	The address family to which the IGP routes will be exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 2048.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines



NOTE

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of an IGP export to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of BGP exported OSPF routes to 1000:

```
configure bgp export ospf shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp import-policy

```
configure bgp import-policy [<policy-name> | none]
```

Description

Configures the import policy for BGP.

Syntax Description

policy-name	Specifies the policy.
none	Specifies no policy.

Default

N/A.

Usage Guidelines

Use the none keyword to remove a BGP import policy.

An import policy is used to modify route attributes while adding BGP routes to the IP route table.

Example

The following command configures a policy *imprt_plcy* for BGP:

```
configure bgp import-policy imprt_plcy
```

The following command unconfigures the import policy for BGP:

```
configure bgp import-policy none
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp local-preference

```
configure bgp local-preference <number>
```

Description

Changes the default local preference attribute.

Syntax Description

number	Specifies a value used to advertise this router's degree of preference to other routers within the AS. Range is 0 to 2147483647.

Default

100.

Usage Guidelines

The range is 0 to 2,147,483,647.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command changes the default local preference attribute to 500:

```
configure bgp local-preference 500
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp med

```
configure bgp med [none | <bgp_med>]
```

Description

Configures the metric to be included in the Multi-Exit-Discriminator (MED) path attribute. The MED path attribute is included in route updates sent to external peers if a value is configured.

Syntax Description

none	Specifies not to use a multi-exist-discriminator number.
bgp_med	Specifies a multi-exist-discriminator number. The range is 0-2147483647.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the metric to be included in the MED path attribute:

```
configure bgp med 3
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor dampening

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]} dampening {{half-life <half-life-minutes> {reuse-limit <reuse-limit-number> suppress-limit <suppress-limit-number> max-suppress <max-suppress-minutes>} | policy-filter [<policy-name> | none]}}
```

Description

Configures route flap dampening over BGP peer sessions.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
half-life	Specifies the dampening half life. Range is 1 to 45 minutes.
reuse	Specifies the reuse limit. Range is 1 to 20000.
suppress	Specifies the suppress limit. Range is 1 to 20000.
max-suppress	Specifies the maximum hold down time. Range is 1 to 255 minutes.
policy-filter	Specifies a policy.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route will be used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route will be suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the `policy-filter` option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for BGP neighbors:

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]] no-dampening
```

Example

The following command configures route flap dampening to the BGP neighbor at 192.168.1.22 to the default values:

```
configure bgp neighbor 192.168.1.22 dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor description

```
configure bgp neighbor [all | <remoteaddr>] description {<description>}
```

Description

Configures a description for a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
description	Specifies a string used to describe the neighbor.

Default

The description is a NULL string by default.

Usage Guidelines

Use this command to attach a description to a BGP neighbor. This description will be displayed in the output of the [show bgp neighbor](#) command when you specify the `detail` option, or when you specify a particular neighbor. Enclose the string in double quotes if there are any blank spaces in the string. The maximum length of the string is 56 characters.

If you do not specify the `<description>` parameter, the description will be reset to the default.

Example

The following command configures the description for the BGP neighbor 192.168.1.22 to *Toledo_5*:

```
configure bgp neighbor 192.168.1.22 description Toledo_5
```

History

This command was first available in ExtremeXOS 11.0.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor maximum-prefix

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} maximum-prefix <number> {{threshold <percent>}} {teardown
{holddown-interval <seconds>}} {send-traps}
```

Description

Configures the maximum number of IP prefixes accepted from a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and console), and/or a trap will be sent to the SNMP manager.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established, if the session is torn down due to maximum prefix exceeded. If the hold-down interval is zero or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the peer group, use the following command:

```
configure bgp peer-group maximum-prefix
```

Example

The following command configures the maximum number of IP prefixes accepted from all neighbors to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp neighbor all maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor next-hop-self

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} [next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the outgoing updates to be the address of the BGP connection originating the update.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (lets BGP decide what would be the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp neighbor 172.16.5.25 next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor no-dampening

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} no-dampening
```

Description

Configures no route flap dampening over BGP peer sessions (disables route flap dampening).

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use the following command to enable route flap dampening for BGP neighbors:

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-
multicast]} dampening [{half-life <half-life-minutes> {reuse-limit <reuse-limit-
number> suppress-limit <suppress-limit-number> max-suppress <max-suppress-minutes>} |
policy-filter [<policy-name> | none]}
```

Example

The following command disables route flap dampening to the BGP neighbor at 192.168.1.22:

```
configure bgp neighbor 192.168.1.22 no-dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor password

```
configure bgp neighbor [all | <remoteaddr>] password [none | {encrypted}  
<tcpPassword>]
```

Description

Configures an MD5 secret password for a neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
none	Specifies not to use a password
encrypted	Specifies an encrypted string; do not use.
tcpPassword	Specifies a password string.

Default

N/A.

Usage Guidelines

Disable the BGP neighbor or the BGP protocol before changing the password.

When a password is configured, TCP MD5 authentication is enabled on the TCP connection that is established with the neighbor.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

To change any one of the following parameters you must disable and re-enable the peer session:

- timer
- source-interface
- soft-in-reset
- password

Changing a route reflector client will automatically disable and enable the peer session.

The `encrypted` option is used by the switch when generating a configuration file, and when parsing a switch-generated configuration file. Do not select the encrypted option in the CLI.

Example

The following command configures the password for a neighbor as *Extreme*:

```
configure bgp neighbor 192.168.1.5 password extreme
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor peer-group

```
configure bgp neighbor [all | <remoteaddr>] peer-group [<peer-group-name> | none] {acquire-all}
```

Description

Configures an existing neighbor as the member of a peer group.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
peer-group-name	Specifies a peer group name.
none	Removes the neighbor from the peer group.
acquire-all	Specifies that all parameters should be inherited by the neighbor from the peer group.

Default

By default, remote AS (if configured for the peer group), source-interface, outbound route policy, send-community and next-hop-self settings are inherited.

Usage Guidelines

If `acquire-all` is not specified, only the default parameters are inherited by the neighbor.

When you remove a neighbor from a peer group, it retains the parameter settings of the group. The parameter values are not reset to those the neighbor had before it inherited the peer group values.

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
```

The new neighbor is created as part of the peer group and inherits all of the existing parameters of the peer group. The peer group must have remote AS configured.

Example

The following command configures an existing neighbor as the member of the peer group *outer*:

```
configure bgp neighbor 192.1.1.22 peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor route-policy

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} route-policy [in | out] [none | <policy>]
```

Description

Configures a route map filter for a neighbor.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
policy	Specifies a policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The policy can be installed on the input or output side of the router. The policy is used to modify or filter the NLRI information and the path attributes associated with it when exchanging updates with the neighbor.

Example

The following command configures the route-policy filter for a neighbor based on the policy *nosales*:

```
configure bgp neighbor 192.168.1.22 route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor route-reflector-client

```
configure bgp neighbor [<remoteaddr> | all] [route-reflector-client | no-  
route-reflector-client]
```

Description

Configures a BGP neighbor to be a route reflector client.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
route-reflector-client	Specifies for the BGP neighbor to be a route reflector client.
no-route-reflector-client	Specifies for the BGP neighbor not to be a route reflector client.

Default

N/A.

Usage Guidelines

Another way to overcome the difficulties of creating a fully-meshed AS is to use *route reflectors*. Route reflectors allow a single router to serve as a central routing point for the AS or sub-AS.

Use this command to implicitly define the router to be a route reflector. The neighbor must be in the same AS as the router.

When changing the route reflector status of a peer, the peer will automatically be disabled and re-enabled and a warning message will appear on the console and in the log.

A *cluster* is formed by the route reflector and its client routers. Peer routers that are not part of the cluster must be fully meshed according to the rules of BGP.

Example

The following command configures a BGP neighbor to be a route reflector client:

```
configure bgp neighbor 192.168.1.5 route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor send-community

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} [send-community | dont-send-community]
```

Description

Configures whether the community path attribute associated with a BGP NLRI should be included in the route updates sent to the BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
send-community	Specifies to include the community path attribute.
dont-send-community	Specifies not to include the community path attribute.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

A BGP community is a group of BGP destinations that require common handling. ExtremeXOS supports the following well-known BGP community attributes:

- no-export
- no-advertise
- no-export-subconfed

Example

The following command includes the community path attribute associated with a BGP NLRI in the route updates sent to all BGP neighbors:

```
configure bgp neighbor all send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor shutdown-priority

```
configure bgp neighbor [all | <remoteaddr>] shutdown-priority <number>
```

Description

Configures the shutdown priority for a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
number	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 1024.

Usage Guidelines



NOTE

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of a BGP neighbor to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of the BGP neighbor 10.0.20.1 to 500:

```
configure bgp neighbor 10.0.20.1 shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor soft-reset

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} soft-reset {in | out}
```

Description

Applies the current input or output routing policy to the routing information already exchanged with the neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output policy is determined by the route policy configured for the neighbor on the input and/or output side of the router. This command does not affect the switch configuration.

If both the local BGP neighbor and the neighbor router support the route refresh capability (ExtremeWare does not support this feature), a dynamic soft input reset can be performed. The “soft-reset input” command will trigger the generation of a route refresh message to be sent to the neighbor. As a response to the “Route-Refresh” message, the neighbor will send the entire BGP routing table in updates.

If the “Route-Refresh” capability is not supported by the neighbor (like ExtremeWare), then the user must preconfigure “soft-input-reset”. If “soft-input-reset” is configured, BGP will store all the incoming routes updates from the neighbor. When the user issues the “soft-input-reset” command, the locally stored incoming routes will be reprocessed against the new policy, and will be installed in the BGP route database.

Example

The following command applies the current input routing policy to the routing information already exchanged with the neighbor:

```
configure bgp neighbor 192.168.1.5 soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor source-interface

```
configure bgp neighbor [<remoteaddr> | all] source-interface [any |  
ipaddress <ipAddr>]
```

Description

Changes the BGP source interface for TCP connections.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.
any	Specifies any source interface.
ipAddr	Specifies the IP address of a source interface.

Default

Any.

Usage Guidelines

The source interface IP address must be a valid IP address of any VLAN configured on the switch.

Example

The following command changes the BGP source interface to 10.43.55.10:

```
configure bgp neighbor 192.168.1.5 source-interface ipaddress 10.43.55.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor timer

```
configure bgp neighbor [<remoteaddr> | all] timer keep-alive <keepalive>  
hold-time <holdtime>
```

Description

Configures the BGP neighbor timers.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.
keepalive	Specifies a BGP neighbor timer keepalive time in seconds. The range is 0 to 21,845 seconds.
holdtime	Specifies a BGP neighbor timer hold time in seconds. The range is 0 and 3 to 65,535 seconds.

Default

The default keepalive setting is 60 seconds. The default hold time is 180 seconds.

Usage Guidelines

The BGP neighbor or BGP protocol must be disabled before changing the timer values.

Example

The following command configures the BGP neighbor timers:

```
configure bgp neighbor 192.168.1.5 timer keep-alive 120 hold-time 360
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp neighbor weight

```
configure bgp neighbor [<remoteaddr> | all] weight <weight>
```

Description

Assigns a locally-used weight to a neighbor connection for the route selection algorithm.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.
weight	Specifies a BGP neighbor weight.

Default

By default, the weight is 0.

Usage Guidelines

All routes learned from this peer are assigned the same weight. The route with the highest weight is more preferable when multiple routes are available to the same network. The range is 0 to 65,535.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command assigns a locally used weight of 10 to a neighbor connection:

```
configure bgp neighbor 192.168.1.5 weight 10
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group dampening

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |
ipv4-multicast]} dampening {{half-life <half-life-minutes> {reuse-limit
<reuse-limit-number> supress-limit <suppress-limit-number> max-suppress
<max-suppress-minutes>}} | policy-filter [<policy-name> | none]}
```

Description

Configures route flap dampening for a BGP peer group.

Syntax Description

peer-group-name	Specifies a peer group
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
half-life-minutes	Specifies the dampening half life.
reuse-limit-number	Specifies the reuse limit.
suppress-limit-number	Specifies the suppress limit.
max-suppress-minutes	Specifies the maximum hold down time.
policy-name	Specifies a policy
none	Removes any policy association.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route will be used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route will be suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the `policy-filter` option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for a BGP peer-group:

```
configure bgp peer-group <peer-group-name> no-dampening
```

Example

The following command configures route flap dampening for the BGP peer group *outer*:

```
configure bgp peer-group outer dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group maximum-prefix

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |
ipv4-multicast]} maximum-prefix <number> {{threshold <percent>}} {teardown
{holddown-interval <seconds>}} {send-traps}
```

Description

Configures the maximum number of IP prefixes accepted for all neighbors in the peer group.

Syntax Description

name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and on the console). An SNMP trap can also be sent.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established, if the session has been torn down due to exceeding the max limit. If the hold down interval is 0 or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the neighbors, use the following command:

```
configure bgp neighbor 192.168.1.1 maximum-prefix
```

Example

The following command configures the maximum number of IP prefixes accepted from the peer group *outer* to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp peer-group outer maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group next-hop-self

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |  
ipv4-multicast]} [next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the updates to be the address of the BGP connection originating the update.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (Let the BGP protocol decide the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp peer-group outer next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group no-dampening

```
configure bgp peer-group <peer-group-name> no-dampening
```

Description

Configures no route flap dampening for a BGP peer group (disables route flap dampening).

Syntax Description

peer-group-name	Specifies a BGP peer group.

Default

This feature is disabled by default.

Usage Guidelines

Use the following command to enable route flap dampening for a BGP peer-group:

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-  
multicast]} dampening {{half-life <half-life-minutes> {reuse-limit <reuse-limit-  
number> supress-limit <suppress-limit-number> max-suppress <max-suppress-minutes>}} |  
policy-filter [<policy-name> | none]}
```

Example

The following command disables route flap dampening to the BGP peer group *outer*:

```
configure bgp peer-group outer no-dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group route-reflector-client

```
configure bgp peer-group <peer-group-name> [route-reflector-client | no-  
route-reflector-client]
```

Description

Configures all the peers in a peer group to be a route reflector client.

Syntax Description

peer-group-name	Specifies a peer group.
route-reflector-client	Specifies that all the neighbors in the peer group be a route reflector client.
no-route-reflector-client	Specifies that all the neighbors in the peer group not be a route reflector client.

Default

N/A.

Usage Guidelines

This command implicitly defines this router to be a route reflector.

The peer group must be in the same AS of this router.

Example

The following command configures the peer group *outer* as a route reflector client:

```
configure bgp peer-group outer route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group send-community

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |  
ipv4-multicast]} [send-community | dont-send-community]
```

Description

Configures whether communities should be sent to neighbors as part of route updates.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
send-community	Specifies that communities are sent to neighbors as part of route updates.
dont-send-community	Specifies that communities are not sent to neighbors as part of route updates.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

Example

The following command configures communities to be sent to neighbors as part of route updates:

```
configure bgp peer-group outer send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group password

```
configure bgp peer-group <peer-group-name> password [none | <tcpPassword>]
```

Description

Configures the TCP MD5 secret password for a peer group and all neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
none	Specifies no password.
tcpPassword	Specifies a password.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the password as *Extreme* for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer password extreme
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group remote-AS-number

```
configure bgp peer-group <peer-group-name> remote-AS-number <number>
```

Description

Configures the remote AS number for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
number	Specifies a remote AS number.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the remote AS number for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer remote-AS-number 65001
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group route-policy

```
configure bgp peer-group <peer-group-name> route-policy [in | out] [none |  
<policy>]
```

Description

Configures the policy for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
in	Specifies to install the policy on the input side.
out	Specifies to install the policy on the output side.
none	Specifies to remove the filter.
policy	Specifies a policy.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the route policy for the peer group *outer* and its neighbors using the policy *nosales*:

```
configure bgp peer-group outer route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group soft-reset

```
configure bgp peer-group <peer-group-name> soft-reset {in | out}
```

Description

Applies the current input/output routing policy to the neighbors in the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

N/A.

Usage Guidelines

The input/output routing policy is determined by the route policy configured for the neighbors in the peer group on the input/output side of the router. This command does not affect configuration of the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command applies the current input routing policy to the neighbors in the peer group *outer*:

```
configure bgp peer-group outer soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group source-interface

```
configure bgp peer-group <peer-group-name> source-interface [any |  
ipaddress <ipAddr>]
```

Description

Configures the source interface for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
any	Specifies any source interface.
ipAddr	Specifies an interface.

Default

N/A.

Usage Guidelines

The source interface IP address must be a valid IP address of a VLAN configured on the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the source interface for the peer group *outer* and its neighbors on 10.34.25.10:

```
configure bgp peer-group outer source-interface ipaddress 10.34.25.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group timer

```
configure bgp peer-group <peer-group-name> timer keep-alive <seconds> hold-
time <seconds>
```

Description

Configures the keepalive timer and hold timer values for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
keep-alive <seconds>	Specifies a keepalive time in seconds. Range is 0 to 21845.
hold-time <seconds>	Specifies a hold-time in seconds. Range is 0 and 3 to 65535.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the keepalive timer and hold timer values for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer timer keep-alive 30 hold-time 90
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp peer-group weight

```
configure bgp peer-group <peer-group-name> weight <number>
```

Description

Configures the weight for the peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
number	Specifies a BGP peer group weight. Range is 0 to 65535.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the weight for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer weight 5
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp restart

```
configure bgp restart [none | planned | unplanned | both | aware-only]
```

Description

Configures the router as a graceful BGP restart router.

Syntax Description

none	Do not act as a graceful BGP restart router.
planned	Only act as a graceful BGP restart router for planned restarts.
unplanned	Only act as a graceful BGP restart router for unplanned restarts.
both	Act as a graceful BGP restart router for both planned and unplanned restarts.
aware-only	Only act as a graceful BGP receiver (helper) router.

Default

The default is none; graceful restart is disabled.

Usage Guidelines

This command configures the router as a graceful BGP router. You can decide to configure a router to enter graceful restart for only planned restarts, for only unplanned restarts, or for both. Also, you can decide to configure a router to be a receiver only (which helps a restarting BGP router to perform the graceful restart process), and not to do graceful restarts itself.

This command cannot be used while BGP is enabled globally on the switch.

Example

The following command configures a router to perform graceful BGP restarts only for planned restarts:

```
configure bgp restart planned
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp restart address-family

```
configure bgp restart [add | delete] address-family [ipv4-unicast | ipv4-multicast]
```

Description

Configures the address family used with graceful BGP restart.

Syntax Description

add	Add the address family.
delete	Remove the address family.
ipv4-unicast	IPv4 unicast addresses.
ipv4-multicast	IPv4 multicast addresses.

Default

The default is IPv4 unicast.

Usage Guidelines

This command configures the address family participating in graceful BGP restart. An address family can be added or deleted. By adding an address family, BGP will instruct the local hardware and software to preserve BGP routes of that address family during a graceful restart. The local OPEN message will contain all the added address families.

This command cannot be used while BGP is enabled globally on the switch.



NOTE

For BGP graceful restart to inter-operate with Cisco routers, any restarting routers connected to Cisco routers must be configured with the command, `enable bgp neighbor capability`, in the following form, `enable bgp neighbor <remoteaddr> capability ipv4-unicast`. The command must be executed before BGP is enabled globally on the switch.

Example

The following command configures a router to add IPv4 unicast addresses to graceful BGP restarts:

```
configure bgp restart add address-family ipv4-unicast
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp restart restart-time

```
configure bgp restart restart-time <seconds>
```

Description

Configures the restart time used with graceful BGP restart. This is the maximum time a receiver router will wait for a restarting router to come back up.

Syntax Description

seconds	Specifies the restart time. The range is 1 to 3600 seconds.

Default

The default is 120 seconds

Usage Guidelines

This command configures the restart timer. This timer is started on the receiver router when it detects the neighbor router is restarting (usually when the peer TCP session is reset). At that time, routes from the restarting router are marked as stale, but are preserved in the routing table. The timer is stopped when the restarting BGP neighbor goes to the ESTABLISHED state (it has finished restarting). If the timer expires, the stale routes are deleted.

Example

The following command configures the graceful BGP restart timer:

```
configure bgp restart restart-time 200
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp restart stale-route-time

```
configure bgp restart stale-route-time <seconds>
```

Description

Configures the stale route timer used with graceful BGP restart. This is the maximum time to hold stale paths on receiver routers while its neighbor gracefully restarts.

Syntax Description

seconds	Specifies the stale route time. The range is 1 to 3600 seconds.

Default

The default is 360 seconds

Usage Guidelines

This command configures the stale route timer. This timer is started when the restarting BGP peer goes to the ESTABLISHED state after it restarts. The timer is stopped when the restarting BGP peer sends EOR messages for all address families. When the timer is stopped, or it expires, the stale routes are deleted.

Example

The following command configures the graceful BGP stale route timer:

```
configure bgp restart stale-route-time 400
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp restart update-delay

```
configure bgp restart update-delay <seconds>
```

Description

Configures the update delay timer used with graceful BGP restart. This is the maximum time to delay updating BGP routes to the local IP route table.

Syntax Description

seconds	Specifies the stale route time. The range is 1 to 3600 seconds.

Default

The default is 600 seconds

Usage Guidelines

This command configures the update delay timer. Usually, a restarting router waits to receive EOR messages from all the receiving BGP neighbors before it starts the route update. Otherwise, it will do the route selection when the timer expires.

Example

The following command configures the graceful BGP update delay timer:

```
configure bgp restart update-delay 800
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp routerid

```
configure bgp routerid <router identifier>
```

Description

Changes the router identifier.

Syntax Description

router identifier	Specifies a router identifier in the IPv4 address format.

Default

N/A.

Usage Guidelines

BGP must be disabled before changing the router ID.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest router ID

Example

The following command changes the router ID:

```
configure bgp routerid 192.1.1.13
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

configure bgp soft-reconfiguration

```
configure bgp soft-reconfiguration
```

Description

Immediately applies the route policy associated with the network command, aggregation, import, and redistribution.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not affect the switch configuration.

Example

The following command applies the route policy associated with the network command, aggregation, import, and redistribution:

```
configure bgp soft-reconfiguration
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

create bgp neighbor peer-group

```
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
```

Description

Creates a new neighbor and makes it part of the peer group.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
peer-group-name	Specifies a peer group.
multi-hop	Specifies to allow connections to EBGp peers that are not directly connected.

Default

N/A.

Usage Guidelines

All the parameters of the neighbor are inherited from the peer group. The peer group should have the remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
configure bgp neighbor [all | <remoteaddr>] peer-group [<peer-group-name> | none]
{acquire-all}
```

If you do not specify acquire-all, only the mandatory parameters are inherited from the peer group. If you specify acquire-all, all of the parameters of the peer group are inherited. This command disables the neighbor before adding it to the peer group.

Example

The following command creates a new neighbor and makes it part of the peer group *outer*:

```
create bgp neighbor 192.1.1.22 peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

create bgp neighbor remote-AS-number

```
create bgp neighbor <remoteaddr> remote-AS-number <number> {multi-hop}
```

Description

Creates a new BGP peer.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
number	Specifies a remote AS number. The range is 1 to 65535.
multi-hop	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N/A.

Usage Guidelines

If the AS number is the same as the AS number provided in the `configure bgp as` command, then the peer is consider an IBGP peer, otherwise the neighbor is an EBGP peer. The BGP session to a newly created peer is not started until the `enable bgp neighbor` command is issued.

Example

The following command creates a new BGP peer:

```
create bgp neighbor 192.168.1.17 remote-AS-number 65001
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

create bgp peer-group

```
create bgp peer-group <peer-group-name>
```

Description

Creates a new peer group.

Syntax Description

peer-group-name	Specifies a peer group.

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 512 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-route-policy
- send-community
- next-hop-self

Each BGP peer group is assigned a unique name when the peer group is created.

Example

The following command creates a new peer group named *outer*:

```
create bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

delete bgp neighbor

```
delete bgp neighbor [<remoteaddr> | all]
```

Description

Deletes one or all BGP neighbors.

Syntax Description

remoteaddr	Specifies the IP address of the BGP neighbor to be deleted.
all	Specifies all neighbors.

Default

N/A.

Usage Guidelines

Use this command to delete one or all BGP neighbors.

Example

The following command deletes the specified BGP neighbor:

```
delete bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

delete bgp peer-group

```
delete bgp peer-group <peer-group-name>
```

Description

Deletes a peer group.

Syntax Description

peer-group-name	Specifies a peer group.

Default

N/A.

Usage Guidelines

Use this command to delete a specific BGP peer group.

Example

The following command deletes the peer group named *outer*:

```
delete bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp

```
disable bgp
```

Description

Disables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable BGP on the router.

Example

The following command disables BGP:

```
disable bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp aggregation

```
disable bgp aggregation
```

Description

Disables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Use this command to disable BGP route aggregation.

Example

The following command disables BGP route aggregation:

```
disable bgp aggregation
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp always-compare-med

```
disable bgp always-compare-med
```

Description

Disables BGP from comparing Multi Exit Discriminators (MEDs) for paths from neighbors in different Autonomous Systems (AS).

Syntax Description

This command has no arguments or variables.

Default

ExtremeXOS does not compare MEDs for paths from neighbors in different AS.

Usage Guidelines

The MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED. By default, during the best path selection process, MED comparison is done only among paths from the same AS.

Example

The following command disables MED from being used in comparison among paths from different AS:

```
disable bgp always-compare-med
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp community format

```
disable bgp community format AS-number : number
```

Description

Disables the AS-number:number format of display for communities in the output of show and upload commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Using this command, communities are displayed as a single decimal value.

Example

The following command disables the AS-number:number format of display for communities:

```
disable bgp community format AS-number : number
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp export

```
disable bgp export [direct | ospf | ospf-extern1 | ospf-extern2 | ospf-
inter | ospf-intra | rip | static] {address-family [ipv4-unicast | ipv4-
multicast]}
```

Description

Disables BGP from exporting routes from other protocols to BGP peers.

Syntax Description

direct	Specifies direct routing.
ospf	Specifies OSPF routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-intra	Specifies OSPF-intra routing.
rip	Specifies RIP routing.
static	Specifies static routing.
address-family	The address family to which the IGP routes will be exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discreet configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and then verify the independent operation of each. Then you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF. Similarly for BGP and ISIS, or BGP and RIP.

You can use policies to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Policies can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command disables BGP from exporting routes from the OSPF protocol to BGP peers:

```
disable bgp export ospf
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp neighbor

```
disable bgp neighbor [<remoteaddr> | all]
```

Description

Disables the BGP session.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

After the session has been disabled, all the information in the route information base (RIB) for the neighbor is flushed.

Example

The following command disables the BGP session:

```
disable bgp neighbor 192.1.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp neighbor capability

```
disable bgp neighbor [all | <remoteaddr>] capability [ipv4-unicast | ipv4-  
multicast | route-refresh]
```

Description

This command disables BGP Multiprotocol (MP) and route-refresh capabilities for neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of the BGP neighbor.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines



NOTE

This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command disables BGP Multiprotocol and route-refresh capabilities for one or all neighbors. After the capabilities have been enabled, the BGP neighbor will announce its capabilities to neighbors in an OPEN message.

Example

The following command disables the route-refresh feature for all neighbors:

```
disable bgp neighbor all route-refresh
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp neighbor remove-private-AS-numbers

```
disable bgp neighbor [<remoteaddr> | all] remove-private-AS-numbers
```

Description

Disables the removal of private AS numbers from the AS path in route updates sent to EBGp peers.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the private AS number can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command disables the removal of private AS numbers from the AS path in route updates sent to the EBGp peers:

```
disable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp neighbor soft-in-reset

```
disable bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast |
ipv4-multicast]} soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

This command can be issued only when both BGP and the BGP neighbor is disabled.

Example

The following command disables the soft input reset for the neighbor at 192.168.1.17:

```
disable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp neighbor use-ip-router-alert

```
disable bgp neighbor [all | <remoteaddr>] use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing BGP messages to the specified neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of the BGP neighbor.

Default

Disabled.

Usage Guidelines

The IP router alert option in a BGP message forces intermediate routers to examine the packet very closely and therefore, indirectly, gives greater reliability that a packet will be delivered to its destination.

Example

The following command disables the feature:

```
disable bgp neighbor 192.168.1.17 use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp peer-group

```
disable bgp peer-group <peer-group-name>
```

Description

Disables a BGP peer group and all its BGP neighbors.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the BGP peer group *outer*:

```
disable bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp peer-group capability

```
disable bgp peer-group <peer-group-name> capability [ipv4-unicast | ipv4-multicast | route-refresh]
```

Description

Disables BGP Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

Usage Guidelines



NOTE

This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command disables BGP Multiprotocol and route-refresh capabilities for a peer group. Once the capabilities are enabled, the BGP peer will announce its capabilities to neighbors in an OPEN message

Example

The following command disables the route-refresh feature for the peer group *outer*:

```
disable bgp peer-group outer route-refresh
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp peer-group remove-private-AS-numbers

```
disable bgp peer-group <peer-group-name> remove-private-AS-numbers
```

Description

Disables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the BGP peer group *outer* from removing private AS numbers:

```
disable bgp peer-group outer remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp peer-group soft-in-reset

```
disable bgp peer-group <peer-group-name> {address-family [ipv4-unicast |  
ipv4-multicast]} soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command disables the soft input reset feature:

```
disable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

disable bgp peer-group use-ip-router-alert

```
disable bgp peer-group <peer-group-name> use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing BGP messages to the specified peer group.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the feature for the peer group *outer*:

```
disable bgp peer-group outer use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp

```
enable bgp
```

Description

Enables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the Border Gateway Protocol (BGP) on the router. Before invoking this command, the local AS number and BGP router ID must be configured.

Example

The following command enables BGP:

```
enable bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp aggregation

```
enable bgp aggregation
```

Description

Enables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:

- 1 Enable aggregation using the following command:

```
enable bgp aggregation
```

- 2 Create an aggregate route using the following command:

```
configure bgp add aggregate-address {address-family [ipv4-unicast | ipv4-multicast]} <ipaddress> {as-match | as-set} {summary-only} {advertise-policy <policy>} {attribute-policy <policy>}
```

Example

The following command enables BGP route aggregation:

```
enable bgp aggregation
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp always-compare-med

```
enable bgp always-compare-med
```

Description

Enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems (ASs) in the route selection algorithm.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

MED is only used when comparing paths from the same AS, unless `always-compare-med` is enabled. When this command is issued, MEDs from different AS are used in comparing paths. A MED value of zero is treated as the lowest MED and therefore the most preferred route.

Example

The following command enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems in the route selection algorithm:

```
enable bgp always-compare-med
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp community format

```
enable bgp community format AS-number : number
```

Description

Enables the as-number:number format of display for the communities in the output of `show` and `upload` commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If not enabled, the communities are displayed as a single decimal value.

Example

The following command enables the AS-number:number format of display for communities:

```
enable bgp community format AS-number : number
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp export

```
enable bgp export [direct | ospf | ospf-extern1 | ospf-extern2 | ospf-inter
| ospf-intra | rip | static] {address-family [ipv4-unicast | ipv4-
multicast]} {export-policy <policy-name>}
```

Description

Enables BGP to export routes from other protocols to BGP peers.

Syntax Description

direct	Specifies direct routing.
ospf	Specifies OSPF routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-intra	Specifies OSPF-intra routing.
rip	Specifies RIP routing.
static	Specifies static routing.
address-family	The address family to which the IGP routes will be exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
policy-name	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discreet configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and then verify the independent operation of each. Then, you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF. Similarly for BGP and ISIS, or BGP and RIP.

You can use a policy to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. A policy can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command enables BGP to export routes from the OSPF protocol to BGP peers:

```
enable bgp export ospf
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp neighbor

```
enable bgp neighbor [<remoteaddr> | all]
```

Description

Enables the BGP session. The neighbor must be created before the BGP neighbor session can be enabled.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
```

Example

The following command enables the BGP neighbor session:

```
enable bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp neighbor capability

```
enable bgp neighbor [all | <remoteaddr>] capability [ipv4-unicast | ipv4-
multicast | route-refresh]
```

Description

This command enables BGP Multiprotocol (MP) and route-refresh capabilities for neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines



NOTE

This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command enables BGP Multiprotocol and route-refresh capabilities for one or all neighbors. After the capabilities have been enabled, the BGP neighbor will announce its capabilities to neighbors in an OPEN message.

Example

The following command disables the route-refresh feature for all neighbors:

```
enable bgp neighbor all capability route-refresh
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp neighbor remove-private-AS-numbers

```
enable bgp neighbor [<remoteaddr> | all] remove-private-AS-numbers
```

Description

Enables the removal of private AS numbers from the AS path in route updates sent to EBGp peers.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the routes can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command enables the removal of private AS numbers from the AS path in route updates sent to the EBGp peers:

```
enable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp neighbor soft-in-reset

```
enable bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast |  
ipv4-multicast]} soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

This command can be issued only when both BGP and the BGP neighbor is disabled.

Example

The following command enables the soft recognition feature:

```
enable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp neighbor use-ip-router-alert

```
enable bgp neighbor [all | <remoteaddr>] use-ip-router-alert
```

Description

Enables the router alert IP option in outgoing BGP messages to the specified neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command will force the IP layer of ExtremeXOS to insert the IP Router Alert Option field in all the outbound BGP messages. IP packets with IP Router Alert option in them examined closely by all the intermediate routers in the transit path, thereby causing transmit delays.

Example

The following command enables the feature:

```
enable bgp neighbor 192.168.1.17 use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp peer-group

```
enable bgp peer-group <peer-group-name>
```

Description

Enables a peer group and all the neighbors of a peer group.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

You can use BGP peer groups to group together up to 200 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-asp-path-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command enables the BGP peer group *outer* and all its neighbors:

```
enable bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp peer-group capability

```
enable bgp peer-group <peer-group-name> capability [ipv4-unicast | ipv4-multicast | route-refresh]
```

Description

This command enables BGP Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

Usage Guidelines



NOTE

This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command enables BGP Multiprotocol and route-refresh capabilities for a peer group. After the capabilities have been enabled, the BGP peer will announce its capabilities to neighbors in an OPEN message.

Example

The following command enables the route-refresh feature for the peer group *outer*:

```
enable bgp peer-group outer capability route-refresh
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp peer-group remove-private-AS-numbers

```
enable bgp peer-group <peer-group-name> remove-private-AS-numbers
```

Description

Enables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the BGP peer group *outer* from removing private AS numbers:

```
enable bgp peer-group outer remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp peer-group soft-in-reset

```
enable bgp peer-group <peer-group-name> {address-family [ipv4-unicast |  
ipv4-multicast]} soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command enables the soft input reset feature:

```
enable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

enable bgp peer-group use-ip-router-alert

```
enable bgp peer-group <peer-group-name> use-ip-router-alert
```

Description

Enables the router alert IP option in outgoing BGP messages to the specified peer group.

Syntax Description

peer-group-name	Specifies a peer group.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the feature for the peer group *outer*:

```
enable bgp peer-group outer use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp

```
show bgp
```

Description

Displays BGP configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays information such as AS number, router ID, local preference, sync flag, route reflection, cluster ID, confederation ID, and AS redistributed networks.

Example

The following command displays BGP configuration information:

```
show bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp neighbor

```
show bgp [neighbor {detail} | {neighbor} <remoteaddr>]
```

Description

Displays information about a specified neighbor.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

Use this command to display information about a specific BGP neighbor. If you do not specify a neighbor, information about all neighbors is displayed.

Example

The following command displays information about a specified neighbor:

```
show bgp neighbor 10.10.10.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp neighbor flap-statistics

```
show bgp neighbor <remoteaddr> {address-family [ipv4-unicast | ipv4-
multicast]} flap-statistics {detail}
[all
| as-path <path-expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>
]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}
]
```

Description

Displays information about neighbor route flap dampening statistics.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
flap-statistics	Specifies that only flap-statistics should be displayed (for route flap dampening enabled routes).
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to display information about BGP neighbor route flap dampening.

The option `network any / <netMaskLen>` will display all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` will display all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays information about a specified neighbor's dampened routes:

```
show bgp neighbor 10.10.10.10 flap-statistics all
```

History

This command was first available in ExtremeXOS 10.1.

The `any / <netMaskLen>` options were added in ExtremeXOS 11.0.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp neighbor routes

```
show bgp neighbor <remoteaddr> {address-family [ipv4-unicast | ipv4-
multicast]} [accepted-routes | received-routes | rejected-routes |
suppressed-routes | transmitted-routes] {detail}
[all
| as-path <path-expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>
]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}
]
```

Description

Displays information about specified neighbor routes.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
accepted-routes	Specifies that only accepted routes should be displayed.
received-routes	Specifies that only received routes should be displayed.
rejected-routes	Specifies that only rejected routes should be displayed.
suppressed-routes	Specifies that only suppressed routes should be displayed (for route flap dampening enabled routes).
transmitted-routes	Specifies that only transmitted routes should be displayed.
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to display information about a specific BGP neighbor routes.

The option `network any / <netMaskLen>` will display all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` will display all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays information about a specified neighbor's received routes:

```
show bgp neighbor 10.10.10.10 received-routes all
```

History

This command was first available in ExtremeXOS 10.1.

The `any / <netMaskLen>` options were added in ExtremeXOS 11.0.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp peer-group

```
show bgp peer-group {detail | <peer-group-name> {detail}}
```

Description

Displays the peer groups configured in the system.

Syntax Description

detail	Specifies to display the information in detailed format.
peer-group-name	Specifies a peer group.

Default

N/A.

Usage Guidelines

If the `detail` keyword is specified then the parameters of the neighbors in the peer group, which are different from the ones that are configured in the peer group, will be displayed.

If no peer group name is specified, all the peer group information will be displayed.

Example

The following command displays the peer groups configured in the system:

```
show bgp peer-group detail
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp routes

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast]} {detail}
[all
| as-path <path-expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>
]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}
]
```

Description

Displays the BGP route information base (RIB).

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The option `network any / <netMaskLen>` will display all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` will display all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays the BGP route information base (RIB):

```
show bgp routes all
```

History

This command was first available in ExtremeXOS 10.1.

The `any / <netMaskLen>` options were added in ExtremeXOS 11.0.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp routes summary

```
show bgp routes summary {address-family [ipv4-unicast | ipv4-multicast]}
```

Description

Displays a summary the BGP route information base (RIB).

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

None.

Example

The following command displays a summary of the BGP route information base (RIB) for IPv4 multicast:

```
show bgp routes summary address-family ipv4-multicast
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

show bgp memory

```
show bgp memory {detail | <memoryType>}
```

Description

Displays BGP specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays BGP specific memory for all types:

```
show bgp memory detail
```

History

This command was first available in ExtremeXOS 10.1.

This command required an Advanced Core license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms with a Core or an Advanced Core license.

IP multicast routing is a function that allows a single IP host to send a packet to a group of IP hosts. This group of hosts can include devices that reside on or outside the local network, or within or across a routing domain.

IP multicast routing consists of the following functions:

- A router that can forward IP multicast packets
- A router-to-router multicast protocol [for example, Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM)]
- A method for the IP host to communicate its multicast group membership to a router [for example, Internet Group Management Protocol (IGMP)]

**NOTE**

You must configure IP unicast routing before you configure IP multicast routing.

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, a single IP host responds to the query, and group registration is maintained.

IGMPv2 is enabled by default on the switch, and beginning in 11.2, ExtremeXOS also supports IGMPv3. However, the switch can be configured to disable the generation of periodic IGMP query packets. IGMP should be enabled when the switch is configured to perform IP unicast or IP multicast routing.

Protocol Independent Multicast (PIM) is a multicast routing protocol with no inherent route exchange mechanism. The switch supports dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

PIM Dense Mode

Protocol Independent Multicast-Dense Mode (PIM-DM) is a multicast routing protocol that is similar to DVMRP. PIM-DM routers perform reverse path forwarding (RPF). However, instead of exchanging its own unicast route tables for the RPF lookup, PIM-DM uses the existing unicast route table for the RPF check. As a result, PIM-DM requires less system memory than DVMRP.

PIM-DM is a broadcast and prune protocol. Using PIM-DM, multicast routes are pruned and grafted in a similar way as DVMRP.

PIM Sparse Mode (PIM-SM)

Unlike PIM-DM, PIM-SM is an explicit join and prune protocol, and it supports shared trees as well as shortest path trees (SPTs). The routers must explicitly join the group(s) in which they are interested in becoming a member, which is beneficial for large networks that have group members who are sparsely distributed.

Using PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. By default, the RP is selected dynamically (but not automatically). You can also define a static RP in your network, using the following command:

```
configure pim crp static <ip_address>
```

When a router has a multicast packet to distribute, it encapsulates the packet in a unicast message and sends it to the RP. The RP decapsulates the multicast packet and distributes it among all member routers.

When a router determines that the multicast rate from a particular group has exceeded a configured threshold, that router can send an explicit join to the originating router. When this occurs, the receiving router gets the multicast directly from the sending router, and bypasses the RP.



NOTE

You can run either PIM-DM or PIM-SM per VLAN.

PIM-SSM

PIM-SM works well in many-to-many multicasting situations. For example, in video conferencing, each participating site multicasts a stream that is sent to all the other participating sites. However, PIM-SM is overly complex for one-to-many multicast situations, such as multimedia content distribution or streaming stock quotes. In these and similar applications, the listener is silent and can know the source of the multicast in advance, or can obtain it. In these situations, there is no need to join an RP, as the join request can be made directly towards the source.

PIM Source Specific Multicast (PIM-SSM) is a special case of PIM-SM, in which a host will explicitly send a request to receive a stream from a specific source, rather than from any source. The host must use IGMPv3 for PIM-SSM, because the ability to request a stream from a specific source first became available with IGMPv3. The PIM-SSM capable router interprets the IGMPv3 message to initiate a PIM-SM join towards the source.

PIM-SSM has the following advantages:

- No overhead of switching to the source-specific tree and waiting for the first packet to arrive
- No need to learn and maintain an RP
- Fewer states to maintain on each router
- No need for the complex register mechanism from the source to the RP
- Better security, as stream is forwarded from sources known in advance

PIM-SSM has the following requirement:

- Any host that participates in PIM-SSM must use IGMPv3

PIM-SSM is designed as a subset of PIM-SM and all messages are compliant with PIM-SM. PIM-SSM and PIM-SM can coexist in a PIM network; only the last hop router needs to be configured for PIM-SSM if both source and receivers are present all the time. However, to avoid any join delay, it is recommended that you enable all routers along the (s,g) path for PIM-SSM.

PIM Mode Interoperation

An Extreme Networks switch can function as a PIM multicast border router (PMBR). A PMBR integrates PIM-SM and PIM-DM traffic.

When forwarding PIM-DM traffic into a PIM-SM network, the PMBR notifies the RP that the PIM-DM network exists. The PMBR forwards PIM-DM multicast packets to the RP, which in turn forwards the packets to those routers that have joined the multicast group in the PIM-SM domain.

The PMBR also forwards PIM-SM traffic to a PIM-DM network. The PMBR sends a join message to the RP and the PMBR floods traffic from the RP into the PIM-DM network.

No commands are needed to enable PMBR mode interoperation. PMBR mode translation is automatically enabled when a dense mode interface and a sparse mode interface are enabled on the same switch.

Multicast VLAN Registration

Multicast VLAN Registration (MVR) is designed to support distributing multicast streams for IPTV to subscribers over a layer 2 network. In a standard layer 2 network, a multicast stream received on a VLAN is not forwarded to another VLAN. The streams are confined to the layer 2 broadcast domain. In an IGMP snooping environment, streams are forwarded only to interested hosts on a VLAN. For inter-VLAN forwarding (routing) a multicast routing protocol, such as PIM/DVMRP must be deployed.

MVR breaks this basic rule, so that a stream received over layer 2 VLANs is forwarded to another VLAN, eliminating the need for a layer 3 routing protocol. It simplifies the multicast stream distribution and is a better solution for IPTV-like services. With MVR, a multicast stream is forwarded to all VLANs containing interested hosts.

clear igmp group

```
clear igmp group {<grpipaddress>} {{vlan} <name>}
```

Description

Removes one or all IGMP groups.

Syntax Description

name	Specifies a VLAN name.
grpipaddress	Specifies the group IP address.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove learned IGMP group entries instantly.

Example

The following command clears all IGMP groups from VLAN *accounting*:

```
clear igmp group accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear igmp snooping

```
clear igmp snooping {{vlan} <name>}
```

Description

Removes one or all IGMP snooping entries.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove IGMP snooping entries instantly. However, removing an IGMP snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The static and dynamic IGMP snooping entries will be removed, then recreated upon the next general query. The static router entry is removed and recreated immediately.

This command clears both the IGMPv2 and IGMPv3 snooping entries.

Example

The following command clears IGMP snooping from VLAN *accounting*:

```
clear igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear pim cache

```
clear pim cache {<group_addr> {<source_addr>}}
```

Description

Resets the IP multicast cache table.

Syntax Description

group_addr	Specifies a group address.
source_addr	Specifies a source IP address.

Default

If no options are specified, all IP multicast cache entries are flushed.

Usage Guidelines

This command can be used by network operators to manually remove IPMC software and hardware forwarding cache entries instantly. If the stream is available, caches will be re-created, otherwise caches are removed permanently. This command can disrupt the normal forwarding of multicast traffic.

Example

The following command resets the IP multicast table for group 224.1.2.3:

```
clear pim cache 224.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp

```
configure igmp <query_interval> <query_response_interval>
<last_member_query_interval> {<robustness>}
```

Description

Configures the Internet Group Management Protocol (IGMP) timers.

Syntax Description

query_interval	Specifies the interval (in seconds) between general queries.
query_response_interval	Specifies the maximum query response time (in seconds).
last_member_query_interval	Specifies the maximum group-specific query response time (in seconds).
robustness	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2236. Specify the following:

- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7. This parameter allows tuning for the expected packet loss on a link. If a link is expected to have packet loss, this parameter can be increased.

Example

The following command configures the IGMP timers:

```
configure igmp 100 5 1 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping flood-list

```
configure igmp snooping flood-list [<policy> | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

policy	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise will be fast path forwarded according to IGMP and/or layer 3 multicast protocol.

A policy file is a text file with the extension, .pol. It can be created or edited with any text editor. The specified policy file <policy file> should contain a list of addresses which will determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the <policy file> in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IP address into the policy file, a 32-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing certain streams as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
#
# This is a template for IGMP Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch the rest of the file!!!!
entry igmpFlood {
    if match any {
#----- Start of group addresses -----
        nlri 234.1.1.1/32;
        nlri 239.1.1.1/32;
#----- end of group addresses -----
    } then {
        permit;
    }
}
```

```

}

entry catch_all {
    if {

        } then {
            deny;
        }
    }
}

```

**NOTE**

The switch will not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to IGMP, PIM, or DVMRP) so it should be used with caution.

Slow path flooding will be done within the L2 VLAN only.

Use the `none` option to effectively disable slow path flooding.

You can use the `show igmp` command to see the configuration of slow path flooding.

Example

The following command configures the multicast data stream specified in *access1* for slow path flooding:

```
configure igmp snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure igmp snooping flood-list none
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping leave-timeout

```
configure igmp snooping leave-timeout <leave_timeout_ms>
```

Description

Configures the IGMP snooping leave timeout.

Syntax Description

leave_timeout_ms	Specifies an IGMP leave timeout value in milliseconds.

Default

1000 ms.

Usage Guidelines

The range is 0 - 10000 ms (10 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

The specified time is the maximum leave timeout value. The switch can leave sooner if an IGMP leave message is received before the timeout occurs.

Example

The following command configures the IGMP snooping leave timeout:

```
configure igmp snooping leave-timeout 10000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping timer

```
configure igmp snooping timer <router_timeout> <host_timeout>
```

Description

Configures the IGMP snooping timers.

Syntax Description

router_timeout	Specifies the time in seconds before removing a router snooping entry.
host_timeout	Specifies the time in seconds before removing a host's group snooping entry.

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- router timeout—The maximum time, in seconds, that a router snooping entry can remain in the IGMP snooping table without receiving a router report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- host timeout—The maximum time, in seconds, that a group snooping entry can remain in the IGMP snooping table without receiving a group report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.

IGMP snooping expects at least one device on every VLAN to periodically generate IGMP query messages. Without an IGMP querier, the switch eventually stops forwarding IP multicast packets to any port, because the IGMP snooping entries will time out, based on the value specified in `host_timeout` or `router_timeout`.

Example

The following command configures the IGMP snooping timers:

```
configure igmp snooping timer 600 600
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping vlan ports add static group

```
configure igmp snooping {vlan} <vlanname> ports <portlist> add static  
group <ip address>
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.
ip address	Specifies the multicast group IP address.

Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group will be forwarded to that port.

This command is for IGMPv2 only.

The switch sends proxy IGMP messages in place of those generated by a real host. The proxy messages use the VLAN IP address for source address of the messages. If the VLAN has no IP address assigned, the proxy IGMP message will use 0.0.0.0 as the source IP address.

The multicast group should be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x/24).

If the ports also have an IGMP filter configured, the filter entries take precedence. IGMP filters are configured using the command:

```
configure igmp snooping vlan <vlan name> ports <portlist> filter <policy file>
```

Example

The following command configures a static IGMP entry so the multicast group 224.34.15.37 will be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static group 224.34.15.37
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping vlan ports delete static group

```
configure igmp snooping {vlan} <vlanname> ports <portlist> delete static  
group [<ip_address> | all]
```

Description

Removes the port configuration that causes multicast group traffic to be forwarded, even if no IGMP leaves have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.
ip address	Specifies the multicast group IP address.
all	Delete all the static groups.

Default

None.

Usage Guidelines

Use this command to remove a static group entry created by the following command:

```
configure igmp snooping vlan <vlan name> ports <portlist> add static  
group <ip address>
```

Example

The following command removes a static IGMP entry that forwards the multicast group 224.34.15.37 to the VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static group 224.34.15.37
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping vlan ports add static router

```
configure igmp snooping {vlan} <vlanname> ports <portlist> add static  
router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic will be forwarded to those ports.

Example

The following command configures a static IGMP entry so all multicast groups will be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping vlan ports delete static router

```
configure igmp snooping vlan <vlanname> ports <portlist> delete static router
```

Description

Removes the configuration that causes VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove an entry created by the following command:

```
configure igmp snooping vlan <vlanname> ports <portlist> add static router
```

Example

The following command removes the static IGMP entry that caused all multicast groups to be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure igmp snooping vlan ports filter

```
configure igmp snooping vlan <vlanname> ports <portlist> filter [<policy> |
none]
```

Description

Configures an IGMP snooping policy file filter on VLAN ports.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
policy	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the class-D addresses of the multicast groups that you wish to block.

To remove IGMP snooping filtering from a port, use the `none` keyword version of the command.

Use the following template to create a snooping filter policy file:

```
#
# Add your group addresses between "Start" and "end"
# Do not touch the rest of the file!!!!

entry igmpFilter {
    if match any {
#----- Start of group addresses -----
        nlri 239.11.0.0/16;
        nlri 239.10.10.4/32;
#----- end of group addresses -----
    } then {
        deny;
    }
}

entry catch_all {
    if {
        } then {
            permit;
        }
    }
}
```

Example

The following command configures the policy file *ap_multicast* to filter multicast packets forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 filter ap_multicast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ipmroute add

```
configure ipmroute add [<source-net>/<mask-len> | <source-net> <mask>]
{{protocol} <protocol>} <rpf-address> {<metric>} {vr <vr-name>}
```

Description

Add a static multicast route to the routing table.

Syntax Description

source-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
protocol	Unicast routing protocol that will be used for route learning.
rpf-address	Next hop through which the multicast source can be reached.
metric	Specifies a cost metric.
vr-name	Specifies the virtual router to which the route is added.

Default

The following defaults apply:

- metric—zero
- vr-name—VR of the current CLI context
- protocol—any

Usage Guidelines

This command allows you to statically configure where multicast sources are located (even though the unicast routing table has different entries). It allows you to configure a multicast static route in such a way as to have non-congruent topology for Unicast and Multicast topology and traffic.

Example

The following command configures a multicast static route for all multicast sources within network subnet 192.168.0.0/16. Those sources are reachable through the gateway 192.75.0.91.

```
configure ipmroute add 192.168.0.0/16 192.75.0.91
```

The following example configures multicast static route for all sources via a single gateway with a metric of 100:

```
configure ipmroute add 0.0.0.0/0 192.75.0.91 100
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure ipmroute delete

```
configure ipmroute delete [<source-net>/<mask-len> | <source-net> <mask>]
{{protocol} <protocol>} <rpf-address> {vr <vr-name>}
```

Description

Deletes a static multicast address from the routing table.

Syntax Description

isource-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
protocol	Unicast routing protocol that will be used for route learning.
rpf-address	Next hop through which the multicast source can be reached.
vr-name	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to delete an existing multicast static route. It allows you to configure congruent topology for unicast and multicast packets and traffic.

Example

The following command deletes a multicast static route:

```
configure ipmroute delete 192.168.0.0/16 192.75.0.91
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure mvr add receiver

```
configure mvr vlan <vlan-name> add receiver port <port-list>
```

Description

Configures a port to receive MVR multicast streams.

Syntax Description

vlan-name	Specifies a VLAN name.
port-list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to add a group of virtual ports for multicast forwarding through MVR. By default, some ports on non-MVR VLANs (router ports, primary and secondary EAPS ports), are excluded from the MVR cache egress list. This command is used to override these rules, so that if valid IGMP memberships are received, or a router is detected, streams are forwarded out on the ports.

Example

The following command adds the ports 1:1 and 1:2 of VLAN *v1* to MVR for forwarding:

```
configure mvr vlan v1 add receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mvr add vlan

```
configure mvr add vlan <vlan-name>
```

Description

Configures a VLAN as an MVR VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Configures MVR on the specified VLAN. When a multicast stream in the specified MVR address range is received on the VLAN, it is leaked to all other VLAN ports where the corresponding IGMP join message is received. By default, the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR. To change the MVR address range, use the following command:

```
configure mvr vlan <vlan-name> mvr-address {<policy-name> | none}
```

Example

The following command configures VLAN *v1* as an MVR VLAN:

```
configure mvr add vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mvr delete receiver

```
configure mvr vlan <vlan-name> delete receiver port <port-list>
```

Description

Configures a port not to receive MVR multicast streams.

Syntax Description

vlan-name	Specifies a VLAN name.
port-list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to delete a group of virtual ports for multicast forwarding through MVR. After using this command, the ports will revert to the default forwarding rules.

Example

The following command deletes the ports 1:1 and 1:2 of VLAN *v1* to MVR for forwarding:

```
configure mvr vlan v1 delete receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mvr delete vlan

```
configure mvr delete vlan <vlan-name>
```

Description

Deletes a VLAN from MVR.

Syntax Description

vlan-name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Removes MVR from the specified VLAN.

Example

The following command configures VLAN *v1* as a non-MVR VLAN:

```
configure mvr delete vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mvr mvr-address

```
configure mvr vlan <vlan-name> mvr-address {<policy-name> | none}
```

Description

Configures the MVR address range on a VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.
policy-name	Specifies a policy file.

Default

The default address range is 224.0.0.0/4 (all multicast addresses), but excluding 224.0.0.0/24 (the multicast control range).

Usage Guidelines

If no policy file is specified (the `none` option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR.

MVR must first be configured on the VLAN before using this command.

If the policy is later refreshed, groups denied and newly allowed groups in the policy will be flushed from fast path forwarding. This will allow synching existing channels with the new policy, without disturbing existing channels.

The following is a sample policy file `mvrpol.pol`. This policy configures 236.1.1.0/24 as the MVR address range. Any address outside this range will have the standard switching behavior on an MVR VLAN.

```
Entry extremel {
    if match any {
        nlri 236.1.1.0/24 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command configures the MVR address range specified in the policy file `mvrpol.pol` for the VLAN `v1`:

```
configure mvr vlan v1 mvr-address mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mvr static group

```
configure mvr vlan <vlan-name> static group {<policy-name> | none}
```

Description

Configures the MVR static group address range on a VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.
policy-name	Specifies a policy file.

Default

By default, all the MVR group addresses work in static mode.

Usage Guidelines

If no policy file is specified (the `none` option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24, is used for static groups in MVR.

MVR must first be configured on the VLAN before using this command.

The following is a sample policy file `mvrpol.pol`. This policy configures 236.1.1.0/24 as the MVR static group address range. Any MVR addresses outside this range will be dynamically registered through IGMP. An MVR VLAN will only proxy join for addresses that are not in the static group. If you want all the multicast groups to be dynamic, use a policy file with this command that denies all multicast addresses.

```
Entry extremel {
    if match any {
        nlri 236.1.1.0/24 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command configures the MVR static group address range specified in the policy file `mvrpol.pol` for the VLAN `v1`:

```
configure mvr vlan v1 static group mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure pim add vlan

```
configure pim add vlan [<vlan-name> | all] {dense | sparse} {passive}
```

Description

Configures an IP interface for PIM.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
dense	Specifies PIM dense mode (PIM-DM).
sparse	Specifies PIM sparse mode (PIM-SM).
passive	Specifies a passive interface.

Default

Dense.

Usage Guidelines

When an IP interface is created, per-interface PIM configuration is disabled by default.

The switch supports both dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

Passive interfaces are host only interfaces that allow a multicast stream from other VLANs to be forwarded to edge hosts. Since they do not peer with other PIM routers, do not connect a multicast router to a passive interface.

In order for the interface to participate in PIM, PIM must be enabled on the switch using the following command:

```
enable pim
```

Example

The following command enables PIM-DM multicast routing on VLAN *accounting*:

```
configure pim add vlan accounting dense
```

History

This command was first available in ExtremeXOS 10.1.

The `passive` parameter was added in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure pim cbsr

```
configure pim cbsr [{vlan} <vlan_name> {<priority [0-254]} | none]
```

Description

Configures a candidate bootstrap router for PIM sparse-mode operation.

Syntax Description

vlan_name	Specifies a VLAN name.
priority	Specifies a priority setting. The range is 0 - 254.
none	Specifies to delete a CBSR.

Default

The default setting for priority is 0, and indicates the lowest priority.

Usage Guidelines

The VLAN specified for CBSR must have PIM enabled for it to take effect. After PIM is enabled, CBSRs advertise themselves in the PIM domain. A BSR is elected among all the candidates based on CBSR priority. To break the tie among routers with the same priority setting, the router with the numerically higher IP address is chosen.

Example

The following command configures a candidate bootstrap router on the VLAN *accounting*:

```
configure pim cbsr vlan accounting 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim crp static

```
configure pim crp static <ip_address> [none | <policy>] {<priority> [0-254]}
```

Description

Configures a rendezvous point and its associated groups statically, for PIM sparse mode operation.

Syntax Description

ip_address	Specifies a static CRP address.
none	Deletes the static rendezvous point.
policy	Specifies a policy file name.
priority	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0, which indicates highest priority.

Usage Guidelines

In PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. If you use a static RP, all switches in your network must be configured with the same RP address for the same group (range).

The policy file contains a list of multicast group addresses served by this RP.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration will look a little different from a policy used for other purposes.

If routers have different group to RP mappings, due to misconfiguration of the static RP (or any other reason), traffic will be disrupted.

Example

The following command statically configures an RP and its associated groups defined in policy file *rp-list*:

```
configure pim crp static 10.0.3.1 rp-list
```

The following is a sample policy file:

```
entry extremel {
    if match any { }
    then { nlri 224.0.0.0/4 ;
          nlri 239.255.0.0/24 ;
          nlri 232.0.0.0/8 ;
          nlri 238.1.0.0/16 ;
```

```
        nlri 232.232.0.0/20 ;  
    }  
}
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim crp timer

```
configure pim crp timer <crp_adv_interval>
```

Description

Configures the candidate rendezvous point advertising interval in PIM sparse mode operation.

Syntax Description

crp_adv_interval	Specifies a candidate rendezvous point advertising interval in seconds. The range is 1 to 1,717,986,918.

Default

The default is 60 seconds.

Usage Guidelines

Increasing this time will result in increased convergence time for CRP information to the PIM routers.

Example

The following command configures the candidate rendezvous point advertising interval to 120 seconds:

```
configure pim crp timer 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim crp vlan

```
configure pim crp vlan <vlan_name> [none | <policy>] {<priority>}
```

Description

Configures the dynamic candidate rendezvous (CRP) point for PIM sparse-mode operation.

Syntax Description

vlan_name	Specifies a VLAN name.
none	Specifies no policy file.
policy	Specifies a policy file name.
priority	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0 and indicates the highest priority.

Usage Guidelines

The policy file contains the list of multicast group addresses serviced by this RP. This set of group addresses are advertised as candidate RPs. Each router then elects the common RP for a group address based on a common algorithm. This group to RP mapping should be consistent on all routers.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration will look a little different from a policy used for other purposes. The following is a sample policy file which will configure the CRP for the address ranges 239.0.0.0/24 and 232.144.27.0/24:

```
entry extremel {
    if match any {
    }
    then {
        nlri 239.0.0.0/24 ;
        nlri 232.144.27.0/24 ;
    }
}
```

The VLAN specified for a CRP must have PIM configured.

To delete a CRP, use the keyword `none` as the access policy.

Example

The following command configures the candidate rendezvous point for PIM sparse-mode operation on the VLAN `HQ_10_0_3` with the policy `rp-list` and priority set to 30:

```
configure pim crp HQ_10_0_3 rp-list 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim delete vlan

```
configure pim delete vlan [<vlanname> | all]
```

Description

Disables PIM on a router interface.

Syntax Description

vlan name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables PIM on VLAN *accounting*:

```
configure pim delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim register-rate-limit-interval

```
configure pim register-rate-limit-interval <interval>
```

Description

Configures the initial PIM-SM periodic register rate.

Syntax Description

interval	Specifies an interval time in seconds. Range is 0 - 60. Default is 0.

Default

Default is 0.

Usage Guidelines

Configuring a non-zero interval time can reduce the CPU load on the first hop switch, in case register stop messages are not received normally.

When a non-zero value is configured, the first hop switch sends a few register messages and then waits for a corresponding register stop from the RP for <time> seconds. The process is repeated until the register stop is received. This command should be used when the (S,G) tree between the first hop router and the RP is not converging quickly.

The default value is zero in default mode, the switch sends continuous register messages until the register stop is received.

Example

The following command configures the initial PIM register rate limit interval:

```
configure pim register-rate-limit-interval 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim register-suppress-interval register-probe-interval

```
configure pim register-suppress-interval <reg-interval> register-probe-
interval <probe_interval>
```

Description

Configures an interval for periodically sending null-registers.

Syntax Description

reg-interval	Specifies an interval time in seconds. Range is 30 - 200 seconds. Default is 60.
probe-interval	Specifies an interval time in seconds. Default is 5.

Default

The following defaults apply:

- register-suppress-interval—60
- register-probe-interval—5

Usage Guidelines

The register-probe-interval time should be set less than the register-suppress-interval time. By default, a null register is sent every 55 seconds (*register-suppress-interval* - *register-probe-interval*). A response to the null register is expected within register probe interval. By specifying a larger interval, a CPU peak load can be avoided because the null-registers are generated less frequently. The register probe time should be less than half of the register suppress time, for best results.

Example

The following command configures the register suppress interval and register probe time:

```
configure pim register-suppress-interval 90 register-probe time 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim register-checksum-to

```
configure pim register-checksum-to [include-data | exclude-data]
```

Description

Configures the checksum computation to either include data (for compatibility with Cisco Systems products) or to exclude data (for RFC-compliant operation), in the register message.

Syntax Description

include-data	Specifies to include data.
exclude-data	Specifies to exclude data.

Default

Include data.

Usage Guidelines

None.

Example

The following command configures the checksum mode to include data for compatibility with Cisco Systems products:

```
configure pim register-checksum-to include-data
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim spt-threshold

```
configure pim spt-threshold <leaf-threshold> {<rp_threshold>}
```

Description

Configures the threshold, in kbps, for switching to SPT. On leaf routers, this setting is based on data packets. On the RP, this setting is based on register packets.

Syntax Description

leaf-threshold	Specifies the rate of traffic per (s,g,v) group in kbps for the last hop. Range is 0 - 419403.
rp_threshold	Specifies an RP threshold. Range is 0 - 419403.

Default

The default setting is 0 for both parameters.

Usage Guidelines

For the best performance, use default value of 0.

Example

The following command sets the threshold for switching to SPT:

```
configure pim spt-threshold 4 16
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim ssm range

```
configure pim ssm range [default | policy <policy-name>]
```

Description

Configures the range of multicast addresses for PIM SSM.

Syntax Description

default	Specifies the default address range, 232.0.0.0/8.
policy-name	Specifies a policy that defines the SSM address range.

Default

By default, no SSM range is configured. Using this command with the `default` keyword sets the range to 232.0.0.0/8. To reset the switch to the initial state, use the `unconfigure pim ssm range` command.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the `disable pim` command.

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command. If you wish to change the PIM SSM range, you must first unconfigure the existing range, and then configure the new range.

SSM requires that hosts use IGMPv3 messages to register to receive multicast group packets. When a range is configured for SSM, any IGMPv2 messages for an address in the range will be ignored. Also, any IGMPv3 Exclude messages will be ignored.



NOTE

If a PIM-SSM range is configured, IGMPv2 messages and IGMPv3 exclude messages within the PIM-SSM range are ignored on all IP interfaces, whether or not PIM-SSM is configured on the interfaces.

To specify a range different from the default PIM SSM range, you will need to create a policy file. The match statement of the policy file contains the group addresses to be treated as PIM SSM addresses. For example, to specify the PIM SSM address range as 232.0.0.0/8 and 233.0.0.0/8, you will use the following policy file:

```
Entry extremel {
    if match any {
        nlri 232.0.0.0/8 ;
        nlri 233.0.0.0/8 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command sets the PIM SSM range to 232.0.0.0/8 and 233.0.0.0/8, if the policy file `ssmrange.pol` contains the policy example used above:

```
configure pim ssm range policy ssmrange.pol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure pim timer vlan

```
configure pim timer <hello_interval> <jp_interval> [{vlan} <vlan_name> |
vlan all]
```

Description

Configures the global PIM timers on the specified router interfaces.

Syntax Description

hello_interval	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1 to 65,535 seconds.
jp_interval	Specifies the join/prune interval. The range is 1 to 65,535 seconds.
vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

- hello_interval—30 seconds.
- jp_interval—60 seconds.

Usage Guidelines

These default timers should only be adjusted when excess PIM control packets are observed on the interface.

Example

The following command configures the PIM timers on the VLAN *accounting*:

```
configure pim timer 150 300 vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure pim vlan trusted-gateway

```
configure pim [{vlan} <vlan_name> | vlan all] trusted-gateway [<policy> | none]
```

Description

Configures a trusted neighbor policy.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.
policy	Specifies an policy file name.
none	Specifies no policy file, so all gateways are trusted.

Default

No policy file, so all gateways are trusted.

Usage Guidelines

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. When the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use a policy file to determine trusted PIM router neighbors for the VLAN on the switch running PIM. This is a security feature for the PIM interface.

Example

The following command configures a trusted neighbor policy on the VLAN *backbone*:

```
configure pim vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable igmp

```
disable igmp {vlan <name>}
```

Description

Disables IGMP on a router interface. If no VLAN is specified, IGMP is disabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

IGMP is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP unicast or IP multicast routing.

This command disables IGMPv2 and IGMPv3.

Example

The following command disables IGMP on VLAN *accounting*:

```
disable igmp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable igmp snooping

```
disable igmp snooping {forward-mcrouter-only | with-proxy | vlan <name>}
```

Description

Disables IGMP snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the IGMP snooping proxy.
name	Specifies a VLAN.

Default

IGMP snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, IGMP snooping is disabled only on that VLAN, otherwise IGMP snooping is disabled on all VLANs.

This command applies to both IGMPv2 and IGMPv3.

If the switch is in the forward-mcrouter-only mode, then the command `disable igmp snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the forward-mcrouter-mode, the command `disable igmp snooping forward-mcrouter-only` has no effect.

To change the snooping mode you must disable IP multicast forwarding. Use the command:

```
disable ipmcforwarding
```

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables IGMP snooping on the VLAN *accounting*:

```
disable igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ipmcforwarding

```
disable ipmcforwarding {vlan <name>}
```

Description

Disables IP multicast forwarding on a router interface.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IP multicast forwarding is disabled by default.

IP forwarding must be enabled before enabling IP multicast forwarding.

Disabling IP multicast forwarding disables any Layer 3 multicast routing for the streams coming to the interface.

Example

The following command disables IP multicast forwarding on the VLAN *accounting*:

```
disable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable mvr

```
disable mvr
```

Description

Disables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables MVR on the system:

```
disable mvr
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable pim

```
disable pim
```

Description

Disables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM on the system:

```
disable pim
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable pim ssm vlan

```
disable pim ssm vlan [<vlan_name> | all]
```

Description

Disables PIM SSM on a router interface.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command disables PIM-SSM on the specified layer 3 VLAN.

IGMPv3 include messages for multicast addresses in the SSM range is only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command disables PIM-SSM multicast routing on VLAN *accounting*:

```
disable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable igmp

```
enable igmp {vlan <vlan name>} {IGMPv1 | IGMPv2 | IGMPv3}
```

Description

Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

vlan name	Specifies a VLAN name.
IGMPv1	Specifies the compatibility mode as IGMPv1.
IGMPv2	Specifies the compatibility mode as IGMPv2.
IGMPv3	Specifies the compatibility mode as IGMPv3.

Default

Enabled, set to IGMPv2 compatibility mode.

Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IP hosts respond to the query, and group registration is maintained.

IGMPv2 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP unicast or IP multicast routing.

Example

The following command enables IGMPv2 on the VLAN *accounting*:

```
enable igmp vlan accounting
```

The following command enables IGMPv3 on the VLAN *finance*:

```
enable igmp vlan finance igmpv3
```

History

This command was first available in ExtremeXOS 10.1.

The IGMPv1, IGMPv2, and IGMPv3 options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable igmp snooping

```
enable igmp snooping {forward-mcrouter-only | vlan <name>}
```

Description

Enables IGMP snooping on the switch.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
name	Specifies a VLAN.

Default

Enabled.

Usage Guidelines

IGMP snooping is a layer 2 function of the switch. It does not require multicast routing to be enabled. The feature reduces the flooding of IP multicast traffic. On the VLAN, IGMP snooping optimizes the usage of network bandwidth and prevents multicast traffic from being flooded to parts of the network that do not need it. The switch does not reduce any IP multicast traffic in the local multicast domain (224.0.0.x).

IGMP snooping is enabled by default on the switch. If you are using multicast routing, IGMP snooping can be enabled or disabled. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN.

If a VLAN is specified, IGMP snooping is enabled only on that VLAN, otherwise IGMP snooping is enabled on all VLANs.

This command applies to both IGMPv2 and IGMPv3.

An optional optimization for IGMP snooping is the strict recognition of routers only if the remote devices are running a multicast protocol. Two IGMP snooping modes are supported:

- The `forward-mcrouter-only` mode forwards all multicast traffic to the multicast router (that is, the router running PIM or DVMRP).
- When not in the `forward-mcrouter-only` mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.



NOTE

The forward-mcrouter-only mode for IGMP snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for IGMP snooping are either in the forward-mcrouter-only mode or in the non-forward-mcrouter-only mode, and not a mixture of the two modes.

To change the snooping mode you must disable IP multicast forwarding. To disable IP multicast forwarding, use the command:

```
disable ipmcforwarding {vlan <name>}
```

To change the IGMP snooping mode from the non-forward-mcrouter-only mode to the forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding
enable igmp snooping forward-mcrouter-only
enable ipmcforwarding {vlan <name>}
```

To change the IGMP snooping mode from the forward-mcrouter-only mode to the non-forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding
disable igmp snooping forward-mcrouter-only
enable ipmcforwarding {vlan <name>}
```

Example

The following command enables IGMP snooping on the switch:

```
enable igmp snooping
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable igmp snooping with-proxy

```
enable igmp snooping with-proxy
```

Description

Enables the IGMP snooping proxy. The default setting is enabled.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

This feature can be enabled when IGMPv3 is enabled; however, it is not effective for IGMPv3.

Example

The following command enables the IGMP snooping proxy:

```
enable igmp snooping with-proxy
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ipmcforwarding

```
enable ipmcforwarding {vlan <name>}
```

Description

Enables IP multicast forwarding on an IP interface.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IPMC forwarding is disabled by default.

IP forwarding must be enabled before enabling IPMC forwarding.

Example

The following command enables IPMC forwarding on the VLAN *accounting*:

```
enable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable mvr

```
enable mvr
```

Description

Enables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables MVR on the system:

```
disable mvr
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable pim

```
enable pim
```

Description

Enables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables PIM on the system:

```
enable pim
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable pim ssm vlan

```
enable pim ssm vlan [<vlan_name> | all]
```

Description

Enables PIM SSM on an IP interface.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command enables PIM-SSM on the specified layer 3 VLAN.

PIM-SM must also be configured on the interface for PIM to begin operating (which includes enabling IP multicast forwarding).

IGMPv3 include messages for multicast addresses in the SSM range will only be processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 include messages in the SSM range will not be processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command enables PIM-SSM multicast routing on VLAN *accounting*:

```
enable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show igmp

```
show igmp {vlan} {<vlan name>}
```

Description

This command can be used to display an IGMP-related configuration and group information, per VLAN.

Syntax Description

vlan name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

The output of this command shows:

- The VLAN name.
- The router interface IP address and subnet mask.
- If the interface is active (up), by the letter U.
- If IP forwarding is enabled for the interface, by the letter f.
- If multicast forwarding is enabled, by the letter M.
- If IGMP is enabled, by the letter i.
- If IGMP snooping is enabled, by the letter z.

Example

The following command displays the IGMP configuration:

```
show igmp
```

The following is sample output from this command:

IGMP:

```
Query Interval: 125 sec
Max Response Time: 10 sec
Last Member Query: 1 sec
Robustness: 2
```

IGMP Snooping:

```
Router Timeout: 260 sec
Host Timeout: 260 sec
Igmp Snooping Fast Leave Time: 1000 ms
Igmp Snooping Flag: forward-all-router
Igmp Snooping Flood-list: none
```

Igmp Snooping Proxy: Enable

VLAN	IP Address		Flags	nLRMA	nLeMA	IGMPver
default	0.0.0.0	/ 0	----z	0	0	2
gho	0.0.0.0	/ 0	----z	0	0	2
hguo_fo	0.0.0.0	/ 0	----z	0	0	2
sqa_east	1.1.1.1	/24	-fmiz	3	0	2
vcs1	12.1.1.115	/24	Ufmiz	6	0	2
vcs2	12.1.2.115	/24	Ufmiz	6	0	2
vcs3	12.2.3.115	/24	-fmiz	3	0	2
vcs4	12.2.4.115	/24	Ufmiz	6	1	2
vcs5	12.2.5.115	/24	-fmiz	3	0	2
vcs6	12.2.6.115	/24	-fmiz	3	0	2
vcs7	12.2.7.115	/24	-fmiz	3	0	2
vcs8	12.2.8.115	/24	-fmiz	3	0	2
vhs1	0.0.0.0	/ 0	U---z	0	4	2
vhs2	117.2.2.115	/24	-fmiz	3	0	2
vhs3	117.2.3.115	/24	-fmiz	3	0	2
vhs4	117.2.4.115	/24	-fmiz	3	0	2
vms1	111.1.1.115	/24	Ufmiz	6	7	2

Flags: (E) Interface Enabled, (i) IGMP Enabled
 (m) Multicast Forwarding Enabled
 (nLeMA) Number of Learned Multicast Addressess
 (nLRMA) Number of Locally registered Multicast Addresses
 (U) Interface Up, (z) IGMP Snooping Enabled

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show igmp group

```
show igmp group {{vlan} {<name>} | {<grpipaddress>}} {IGMPv3}
```

Description

Lists the IGMP group membership for the specified VLAN.

Syntax Description

grpipaddress	Specifies a group IP address.
name	Specifies a VLAN name.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

The output of this command shows:

- The multicast group address received.
- The version of the IGMP group.
- The name of the VLAN where the group address is being received.
- The physical port where the group address is being received. If multiple ports within the VLAN have subscribers for the group, all the ports are listed.
- The age since the last IGMP report for this group was received.



NOTE

The show igmp group command output is populated on the router that is the PIM Rendezvous Point.

Example

The following command lists the IGMP group membership:

```
show igmp group
```

The following is sample output from this command:

Group Address	Ver	Vlan	Port	Age
239.2.4.70	2	banana	7	101
224.0.1.24	2	banana	7	107
239.255.255.254	2	banana	7	103

Total: 3

History

This command was first available in ExtremeXOS 10.1.

The IGMPv3 option was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show igmp snooping

```
show igmp snooping {vlan <name> {port <port>}| detail | cache} {IGMPv3}
```

Description

Displays IGMP snooping registration information and a summary of all IGMP timers and states.

Syntax Description

name	Specifies a VLAN name.
detail	Displays the information in detailed format.
cache	Displays the cache setting for IGMP snooping senders.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

The two types of IGMP snooping entry are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

The following command displays IGMP snooping registration information for the switch:

```
show igmp snooping vlan accounting
```

Output from this command looks similar to the following:

Vlan	Vid	Port	#Senders	#Receivers	Router	Enable
default	1		0			Yes
vhs3	4090		0			Yes
vhs4	4089		0			Yes
vcs5	15		0			Yes
vcs6	16		0			Yes
vcs3	4086		0			Yes
vcs4	1014		0			Yes
		5:7		5	No	
		5:9		5	No	

	5:10	5	No	
	5:11	1	No	
	5:12	5	No	
	5:37	5	No	
	5:39	5	No	
	5:41	5	No	
	5:42	5	No	
vcs7	4084	0		Yes
vcs8	4083	0		Yes
vhs2	4082	0		Yes
hguo_fo	200	0		Yes
vcs1	12	8		Yes
	4:16	0	Yes	
vcs2	22	8		Yes
	4:16	0	Yes	
vhs1	1717	14		Yes
	4:32	0	Yes	
vms1	111	2		Yes
	4:10	5	Yes	
gho	4061	0		Yes
sqa_east	4059	0		Yes

The following command displays IGMP snooping registration information on the VLAN *accounting*:

```
show igmp snooping vlan accounting
```

Output from this command looks similar to the following:

```
VLAN accounting (4093) Snooping=Enabled
Port Host          Subscribed  Age
5:10 192.206.38.103 235.1.1.1(s) 0
5:12 192.206.38.103 235.1.1.1(s) 0
5:13 192.206.38.103 235.1.1.1(s) 0
5:14 192.206.38.103 235.1.1.1(s) 0
s = static igmp member
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show igmp snooping vlan filter

```
show igmp snooping {vlan} <name> filter
```

Description

Displays IGMP snooping filters.

Syntax Description

name	Specifies a VLAN name.

Default

None.

Usage Guidelines

Use this command to display IGMP snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters will be displayed.

Example

The following command displays the IGMP snooping filter configured on VLAN *vlan101*:

```
show igmp snooping vlan101 filter
```

The output of the command will be similar to the following:

```
Filter          Port Flags
igmppermit0    5:10 a
```

```
Flags: (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show igmp snooping vlan static

```
show igmp snooping vlan <name> static [group | router]
```

Description

Displays static IGMP snooping entries.

Syntax Description

name	Specifies a VLAN name.

Default

None.

Usage Guidelines

Use this command to display the IGMP snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports will be displayed.

Example

The following command displays the IGMP snooping static groups configured on VLAN *vlan101*:

```
show igmp snooping vlan101 static group
```

The output of the command will be similar to the following:

```
VLAN vlan101 (4094)
  Group      Port  Flags
  239.1.1.2   29    s-
  239.1.1.2   30    s-
  239.1.1.2   31    sa
  239.1.1.2   32    s-
  239.1.1.2   34    s-

Total number of configured static IGMP groups = 5
Flags: (s) Static, (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ipmroute

```
show ipmroute {<source-net>/<mask-len> | <source-net> <mask> | summary} {vr
<vr-name>}
```

Description

Displays the contents of the IP multicast routing table or the route origin priority.

Syntax Description

isource-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
summary	Displays the statistics of multicast static routes.
vr-name	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to view the configured multicast static routes. You can specify the filtering criteria on this CLI to view only the desired route. The multicast static routes are displayed in ascending order of their prefix (same order as `show iproute` displays).

Example

The following example displays a multicast static route from a default virtual router:

```
* (debug) Summit-PC.19 # show ipmroute
```

Destination	Gateway	Mtr	Flags	Protocol	VLAN
Default Route	20.20.20.1	255	UG	ANY	pc4-1
*1.1.0.0/16	20.20.20.1	10	UG	bgp	pc4-1
*11.0.0.0/8	30.30.30.1	12	U-	ANY	pc5-3
11.22.0.0/16	20.20.20.1	10	UG	ANY	pc4-1
*11.22.33.0/24	30.30.30.1	8	U-	ANY	pc5-3
11.22.33.44/32	20.20.20.1	4	UG	ANY	pc4-1
*12.0.0.0/8	20.20.20.1	0	UG	ANY	pc4-1
12.24.0.0/16	30.30.30.1	0	U-	ANY	pc5-3
*12.24.48.96/32	30.30.30.1	2	U-	ospf-extern1	pc5-3
44.66.0.0/16	30.30.30.1	0	U-	ANY	pc5-3

Flags: (*) Active, (G) Gateway, (U) Up

Mask distribution:

1 default routes	2 routes at length 8
4 routes at length 16	1 routes at length 24
2 routes at length 32	

Total number of multicast static routes = 10

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show L2stats

```
show L2stats {vlan <vlan_name>}
```

Description

Displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics).

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics) for the VLAN *accounting*:

```
show L2stats accounting
```



NOTE

You can also enter the command as show l2stats. We use the uppercase letter here to avoid confusion with the numeral 1.

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show mvr

```
show mvr {vlan <vlan_name>}
```

Description

Displays the MVR configuration on the switch.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If a VLAN is specified, information for the VLAN is displayed.

Example

The following command displays the MVR configuration for the VLAN *accounting*:

```
show mvr accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show mvr cache

```
show mvr cache {vlan <vlan_name>}
```

Description

Displays the multicast cache in the MVR range.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If no VLAN is specified, information for all the VLANs is displayed.

Example

The following command displays the multicast cache in the MVR range for the VLAN *accounting*:

```
show mvr cache accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show pim

```
show pim {detail | rp-set {<group_addr>} | vlan <vlan_name>}
```

Description

Displays the PIM configuration and statistics.

Syntax Description

detail	Specifies to display the detailed format.
group_addr	Specifies an IP multicast group, for which the RP is to be displayed.
vlan_name	Specifies a VLAN name.

Default

If no VLAN is specified, the configuration is displayed for all PIM interfaces.

If no multicast group is specified for the `rp-set` option (Rendezvous Point set), all RPs are displayed.

Usage Guidelines

The `detail` version of this command displays the global statistics for PIM, as well as the details of each PIM enabled VLAN.

Examples

The following command displays the global PIM configuration and statistics:

```
show pim
```

The following is sample output from the command:

```
PIM Enabled, Version 2
PIM CRP Enabled on 1 interfaces
BSR state : ELECTED ; BSR Hash Mask Length: 255.255.255.252
Current BSR Info : 40.0.0.1 (Priority 0) expires after 54 sec
Configured BSR Info : 40.0.0.1 (Priority 0) in vlan v_lhr
CRP Adv Interval : 60 sec ; CRP Holdtime: 150
BSR Interval : 60 sec ; BSR Timeout : 130
Cache Timer : 210 sec ; Prune Timer : 210
Assert Timeout : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id : 0x460030fb
Threshold for Last Hop Routers: 0 kbps
Threshold for RP : 0 kbps
Register-Rate-Limit-Interval : Always active
PIM SSM address range : None
Register Checksum to include data
Active Sparse Ckts 2 Dense Ckts 0
Global Packet Statistics (In/Out)
```

```

C-RP-Advs 213 214
Registers 212 0
RegisterStops 0 212
VLAN Cid IP Address Designated Flags Hello J/P Nbrs
Router Int Int
v_bdr 1 30.0.0.2 /24 30.0.0.2 rifms----b 30 60 1
v_lhr 2 40.0.0.1 /24 40.0.0.2 rifmsc---- 30 60 1
Legend: J/P Int: Join/Prune Interval
Flags : r - Router PIM Enabled, i - Interface PIM Enabled, f - Interface,
Forwarding Enabled, m - Interface Multicast Forwarding Enabled,
s - Sparse mode, d - Dense mode, c - CRP enabled,
t - Trusted Gateway configured, n - Multinetted VLAN,
p - Passive Mode, S - Source Specific Multicast, b - Border.

```

The following command displays the detailed PIM configuration and statistics:

```
show pim v_bdr
```

The following is sample output from the command:

```

PIM SPARSE Interface[1] on VLAN v_bdr is enabled and up
IP adr: 30.0.0.2 mask: 255.255.255.0 DR of the net: 30.0.0.2
Passive : No
Hello Interval : 30 sec
Neighbor Time out : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway : none
CRP group List : none with priority 0
Shutdown priority : 1024
Source Specific Multicast : Disabled
Border : Yes
Neighbor IP address Generation Id Expires
30.0.0.1 0x4600309d 79
Packet Statistics (In/Out)
Hellos 427 428 Bootstraps 0 0
Join/Prunes 214 0 Asserts 0 0
Grafts 0 0 GraftAcks 0 0

```

The following command displays the elected, active RP for the group 239.255.255.1:

```
show pim rp-set 239.255.255.1
```

The following is sample output from the command:

Group	Mask	C-RP	Origin	Priority
224.0.0.0	240.0.0.0	10.10.10.2	Bootstrap	0
224.0.0.0	240.0.0.0	124.124.124.124	Bootstrap	0
224.0.0.0	240.0.0.0	124.124.124.124	static	0
239.255.255.0	255.255.255.0	124.124.124.124	Bootstrap	0

Elected RP is 124.124.124.124

History

This command was first available in ExtremeXOS 10.1.

The PIM-SSM information was added in ExtremeXOS 11.4.

Border VLAN information was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show pim cache

```
show pim cache { {detail} {<group_addr> {<source_addr>} } }
```

Description

Displays the IP multicast forwarding cache.

Syntax Description

detail	Specifies to display the information in detailed format.
group_addr	Specifies an IP group address.
source_addr	Specifies an IP source address.

Default

N/A.

Usage Guidelines

Displays the following information:

- IP group address
- IP source address / source mask
- Upstream neighbor (RPF neighbor)
- Interface (VLAN) to upstream neighbor
- Cache expire time
- Egress and prune interface list

When the detail option is specified, the switch displays the egress VLAN list and the pruned VLAN list.

Example

The following command displays the PIM cache entry for group 239.255.255.1:

```
show pim cache 239.255.255.1
```

The following is sample output from this command:

```
Index  Dest Group      Source              InVlan  Origin
[0000] 239.255.255.1  124.124.124.124 (WR) v4      Sparse
Entry timer is not run; UpstNbr: 200.124.124.24
EgressIfList = vbs15(0) (FW) (SM) (I)

[0001] 239.255.255.1  118.5.1.1 (S)      vbs5  Sparse
Expires after 186 secs UpstNbr: 0.0.0.0
RP: 124.124.124.124 via 200.124.124.24 in v4
EgressIfList = vbs15(0) (FW) (SM) (I) , vpim5(170) (FW) (SM) (S)
PrunedIfList = v4(0) (SM)
```

Number of multicast cache = 20

Entry flags :-

R: RP tree. S: Source tree. W: Any source.

Egress/Pruned interface flags :-

SM: Sparse Mode

DM: Dense Mode

Fw: Forwarding

PP: Prune pending

AL: Assert Loser

N: Neighbor present

I: IGMP member present S: (s,g) join received

Z: (*,g) join received Y: (*,*,rp) join received

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure igmp

```
unconfigure igmp
```

Description

Resets all IGMP settings to their default values and clears the IGMP group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all IGMP settings to their default values and clears the IGMP group table:

```
unconfigure igmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure pim

```
unconfigure pim {vlan <vlan_name>}
```

Description

Resets all PIM settings on one or all VLANs to their default values.

Syntax Description

vlan_name	Specifies the VLAN from which PIM is to be unconfigured.

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.

Usage Guidelines

If you unconfigure PIM, you will also unconfigure PIM-SSM, removing the PIM-SSM range.

Example

The following command resets all PIM settings on the VLAN *accounting*:

```
unconfigure pim vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure pim ssm range

```
unconfigure pim ssm range
```

Description

Unconfigures the range of multicast addresses for PIM SSM.

Syntax Description

This command has no arguments or variables.

Default

By default, no SSM range is configured.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the [disable pim](#) command.

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command.

When no range is configured for PIM SSM, the switch will not use PIM SSM for any multicast groups.

Example

The following command removes the PIM SSM range:

```
unconfigure pim ssm range
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

IPv6 multicast is a function that allows a single IPv6 host to send a packet to a group of IPv6 hosts. This group of hosts can include devices that reside on or outside the local network, or within or across a routing domain.

IPv6 multicast routing consists of the following functions:

- A router that can forward IPv6 multicast packets
- A router-to-router multicast protocol [for example, Protocol Independent Multicast (PIM)]
- A method for the IPv6 host to communicate its multicast group membership to a router [for example, Multicast Listener Discovery (MLD)]

**NOTE**

You must configure IPv6 unicast routing before you configure IPv6 multicast routing.

MLD is a protocol used by an IPv6 host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, a single IP host responds to the query, and group registration is maintained.

MLDv1 is currently supported on the switch.

**NOTE**

In the current release of ExtremeXOS (11.4), IPv6 multicast packets are flooded to VLANs that receive the traffic, therefore, MLD snooping is not supported. Also, ExtremeXOS does not currently support any IPv6 multicast routing protocols.

clear mld group

```
clear mld group {<v6grpipaddress>} [{vlan} <name>}
```

Description

Removes one or all MLD groups.

Syntax Description

name	Specifies a VLAN name.
v6grpipaddress	Specifies the group IP address.

Default

N/A.

Usage Guidelines

This command is used to manually remove learned MLD group entries instantly.

Example

The following command clears all MLD groups from VLAN *accounting*:

```
clear mld group accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

clear mld snooping

```
clear mld snooping [{vlan} <name>]
```

Description

Removes one or all MLD snooping entries.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove MLD snooping entries instantly. However, removing an MLD snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The static and dynamic MLD snooping entries will be removed, then recreated upon the next general query. The static router entry is removed and recreated immediately.

Example

The following command clears MLD snooping from VLAN *accounting*:

```
clear mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld

```
configure mld <query_interval> <query_response_interval>
<last_member_query_interval> {<robustness>}
```

Description

Configures the Multicast Listener Discovery (MLD) timers.

Syntax Description

query_interval	Specifies the interval (in seconds) between general queries.
query_response_interval	Specifies the maximum query response time (in seconds).
last_member_query_interval	Specifies the maximum group-specific query response time (in seconds).
robustness	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2710. Specify the following:

- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7.

Example

The following command configures the MLD timers:

```
configure mld 100 5 1 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping vlan ports add static group

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static
group <v6grpipaddress>
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
v6grpipaddress	Specifies the multicast group IPv6 address.

Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group will be forwarded to that port.

The switch sends proxy MLD messages in place of those generated by a real host. The proxy messages use the VLAN IPv6 address for source address of the messages. If the VLAN has no IPv6 address assigned, the proxy MLD message will use 0::0 as the source IP address.



NOTE

In the current implementation, multicast traffic is flooded to the VLAN.

Example

The following command configures a static MLD entry so the multicast group ff02::1:1 will be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static group ff02::1:1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping vlan ports delete static group

```
configure mld snooping {vlan} <vlanname> ports <portlist> delete static
group [all | <v6grpipaddress>]
```

Description

Removes the configuration that causes VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all multicast groups.
v6grpipaddress	Specifies the multicast group IPv6 address.

Default

None.

Usage Guidelines

Use this command to delete a static group from a particular VLAN port.

To add a static group, use the following command:

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static
group <v6grpipaddress>
```

Example

The following command removes a static MLD entry so the multicast group ff02::a:b will not be forwarded to VLAN *marketing* on ports 2:1-2:4, unless an MLD join message is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static group ff02::a:b
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping vlan ports add static router

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic will be forwarded to those ports.

Example

The following command configures a static MLD entry so all multicast groups will be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping vlan ports delete static router

```
configure mld snooping {vlan} <vlanname> ports <portlist> delete static
router
```

Description

Configures VLAN ports to stop forwarding the traffic from all multicast groups, unless MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove the configuration that forwards all multicast groups to the specified VLAN ports.

Example

The following command removes a static MLD entry so all multicast groups will not be forwarded to VLAN *marketing* on ports 2:1-2:4, unless an MLD join is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping flood-list

```
configure mld snooping flood-list [<policy> | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

policy	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise will be fast path forwarded according to MLD and/or layer 3 multicast protocol.

A policy file is a text file with the extension .pol. It can be created or edited with any text editor. The specified policy file <policy file> should contain a list of addresses which will determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the <policy file> in 'permit' mode, that stream will be software flooded and no hardware entry would be installed.

When adding an IPv6 address into the policy file, a 128-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing certain stream as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
#
# This is a template for MLD Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch rest of file!!!!
entry mldFlood {
    if match any {
#----- Start of group addresses -----
        nlri  ff05::100:1/128;
        nlri  ff05::100:15/128;
#----- end of group addresses -----
    } then {
        permit;
    }
}
```

```

}

entry catch_all {
    if {

        } then {
            deny;
        }
    }
}

```

**NOTE**

The switch will not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to MLD or PIM) so it should be used with caution.

Slow path flooding will be done within the L2 VLAN only.

Use the `none` option to effectively disable slow path flooding.

You can use the `show mld` command to see the configuration of slow path flooding.

**NOTE**

This command has no effect in the current release (11.3), since IPv6 multicast traffic floods on all platforms.

Example

The following command configures the multicast data stream specified in *access1* for slow path flooding:

```
configure mld snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure mld snooping flood-list none
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping leave-timeout

```
configure mld snooping leave-timeout <leave_timeout_ms>
```

Description

Configures the MLD snooping leave timeout.

Syntax Description

leave_timeout_ms	Specifies an MLD leave timeout value in milliseconds, upon receiving an MLD done message.

Default

1000 ms.

Usage Guidelines

The range is 0 - 10000 ms (10 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

The specified time is the maximum leave timeout value. The switch could leave sooner if an MLD done message is received before the timeout occurs.

Example

The following command configures the MLD snooping leave timeout:

```
configure mld snooping leave-timeout 10000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mld snooping timer

```
configure mld snooping timer <router_timeout> <host_timeout>
```

Description

Configures the MLD snooping timers.

Syntax Description

router_timeout	Specifies the time in seconds before removing a router snooping entry.
host_timeout	Specifies the time in seconds before removing a host's group snooping entry.

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- router timeout—The maximum time, in seconds, that a router snooping entry can stay without receiving a router report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- host timeout—The maximum time, in seconds, that a group snooping entry can stay without receiving a group report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.

MLD snooping is a layer 2 function of the switch. It does not require multicast routing to be enabled. The feature reduces the flooding of IPv6 multicast traffic. On the VLAN, MLD snooping optimizes the usage of network bandwidth and prevents multicast traffic from being flooded to parts of the network that do not need it. The switch does not reduce any IP multicast traffic in the local multicast domain (FF02::x).

MLD snooping is enabled by default on the switch. MLD snooping expects at least one device on every VLAN to periodically generate MLD query messages. Without an MLD querier, the switch eventually stops forwarding IPv6 multicast packets to any port, because the MLD snooping entries will time out, based on the value specified in `host timeout`.

Example

The following command configures the MLD snooping timers:

```
configure mld snooping timer 600 600
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable mld

```
disable mld {vlan <name>}
```

Description

Disables MLD on a router interface. If no VLAN is specified, MLD is disabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

MLD is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of MLD packets. MLD should be enabled when the switch is configured to perform IPv6 unicast or IPv6 multicast routing.

This command disables all MLD versions.

Example

The following command disables MLD on VLAN *accounting*:

```
disable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable mld snooping

```
disable mld snooping {forward-mcrouter-only | with-proxy | vlan <name>}
```

Description

Disables MLD snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the MLD snooping proxy.
name	Specifies a VLAN.

Default

MLD snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, MLD snooping is disabled only on that VLAN, otherwise MLD snooping is disabled on all VLANs.

If the switch is in the forward-mcrouter-only mode, then the command `disable mld snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the forward-mcrouter-mode, the command `disable mld snooping forward-mcrouter-only` has no effect.

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD done messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables MLD snooping on the VLAN *accounting*:

```
disable mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable mld

```
enable mld {vlan <vlan name>} {MLDv1 | MLDv2}
```

Description

Enables MLD on a router interface. If no VLAN is specified, MLD is enabled on all router interfaces.

Syntax Description

vlan name	Specifies a VLAN name.
MLDv1	Specifies the compatibility mode to MLDv1.
MLDv2	Specifies the compatibility mode to MLDv2.

Default

Enabled, set to MLDv1 compatibility mode.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IPv6 hosts respond to the query, and group registration is maintained.

MLDv1 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of MLD packets.

A VLAN must have an IPv6 address to support MLD.



NOTE

MLDv2 is not supported in ExtremeXOS 11.3.

Example

The following command enables MLDv1 on the VLAN *accounting*:

```
enable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable mld snooping

```
enable mld snooping {forward-mcrouter-only | vlan <name>}
```

Description

Enables MLD snooping on the switch.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
name	Specifies a VLAN.

Default

Enabled.

Usage Guidelines

If a VLAN is specified, MLD snooping is enabled only on that VLAN, otherwise MLD snooping is enabled on all VLANs.

A VLAN must have an IPv6 address to support MLD.

Two MLD snooping modes are supported:

- The `forward-mcrouter-only` mode forwards all multicast traffic to the multicast router (that is, the router running PIM).
- When not in the `forward-mcrouter-only` mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.



NOTE

The forward-mcrouter-only mode for MLD snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for MLD snooping are either in the forward-mcrouter-only mode or in the non-forward-mcrouter-only mode, and not a mixture of the two modes.

To change the MLD snooping mode from the non-forward-mcrouter-only mode to the forward-mcrouter-only mode, use the command:

```
enable mld snooping forward-mcrouter-only
```

To change the MLD snooping mode from the forward-mcrouter-only mode to the non-forward-mcrouter-only mode, use the command:

```
disable mld snooping forward-mcrouter-only
```

Example

The following command enables MLD snooping on the switch:

```
enable mld snooping
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable mld snooping with-proxy

```
enable mld snooping with-proxy
```

Description

Enables the MLD snooping proxy. The default setting is enabled.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

Example

The following command enables the MLD snooping proxy:

```
enable mld snooping with-proxy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show mld

```
show mld {vlan} {<vlan name>}
```

Description

This command can be used to display an MLD-related configuration and group information, per VLAN or for the switch as a whole.

Syntax Description

vlan name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If you do not specify a VLAN, the command displays the switch configuration.

Example

The following command displays the MLD configuration:

```
show mld
```

The following is sample output from this command:

```
MLD:
  Query Interval: 125 sec
  Max Response Time: 10 sec
  Last Member Query: 1 sec
  Robustness: 2

MLD Snooping:
  Router Timeout: 260 sec
  Host Timeout: 260 sec
  MLD Snooping Fast Leave Time: 1000 ms
  MLD Snooping Flag: forward-all-router
  MLD Snooping Flood-list: none
  MLD Snooping Proxy: Enable
```

VLAN	IP Address	Flags	nLRMA	nLeMA	MLDver
Default	::/0	U--iz	0	0	1
loopb	::/0	U--iz	0	0	1
red	fe80::2e0:81ff:fe22:5724/64	U--iz	0	0	2
yellow	::/0	U--iz	0	0	1

Flags: (E) Interface Enabled, (i) MLD Enabled
 (f) Forwarding Enabled, (m) Multicast Forwarding Enabled
 (nLeMA) Number of Learned Multicast Addressess
 (nLRMA) Number of Locally registered Multicast Addresses
 (U) Interface Up, (z) MLD Snooping Enabled

The following command displays the MLD configuration for the VLAN *red*:

```
show mld red
```

The following is sample output from this command:

```
Interface on VLAN red is enabled and up.
  inet6 fe80::2e0:81ff:fe22:5724/64
  Locally registered multicast addresses:

  Learned multicast addresses (Last Querier=fe80::2e0:81ff:fe22:5724):

      s = static igmp member

Flags:
  IP Fwding   NO      IPmc Fwding  NO      MLD YES
  MLD Ver    v1      Snooping   YES
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show mld group

```
show mld group {{vlan} {<name>} | {<v6grpipaddress>}} {MLDv2}
```

Description

Lists the MLD group membership for the specified VLAN.

Syntax Description

^{v6} grpipaddress	Specifies a group IPv6 address.
name	Specifies a VLAN name.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format).

Default

MLDv1.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

Example

The following command lists the MLD group membership for the VLAN *accounting*:

```
show mld group vtest3
```

Output from this command looks similar to the following:

Group Address	Ver	Vlan	Port	Age
ff03::1:1	1	vtest3	4:5	25
ff03::1:2	1	vtest3	4:5	25
ff02::1:ff22:124	1	vtest3	4:45	26
ff05::a:abcd	1	vtest3	4:15	23
ff05::a:abce	1	vtest3	4:15	23
ff02::1:ff22:112	1	vtest3	4:45	26
ff02::1:ff1f:a418	1	vtest3	4:45	26

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show mld snooping

```
show mld snooping {vlan <name> | detail} {MLDv2}
```

Description

Displays MLD snooping registration information and a summary of all MLD timers and states.

Syntax Description

name	Specifies a VLAN name.
detail	Displays the information in detailed format.
cache	Displays the cache setting for MLD snooping senders.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format).

Default

MLDv1.

Usage Guidelines

The two types of MLD snooping entries are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

The following command displays MLD snooping registration information for the VLAN *red*:

```
show mld snooping vlan test3
```

Output from this command looks similar to the following:

```
VLAN vtest3          (133) Snooping=Enabled
  Port   Host                Subscribed                Age
  4:5    fe80::200:ff:fe00:1  ff03::1:1                10
  4:5    fe80::200:ff:fe00:1  ff03::1:2                10
  4:15   fe80::200:ff:fe00:201 ff05::a:abcd              8
  4:15   fe80::200:ff:fe00:201 ff05::a:abce              8
```

4:45	fe80::204:96ff:fe1f:a418	ff02::1:ff1f:a418	11
4:45	fe80::204:96ff:fe1f:a418	ff02::1:ff22:112	11
4:45	fe80::201:30ff:fef9:9b90	ff02::1:ff22:124	11
4:45	fe80::201:30ff:fef9:9b90	All Groups	11

s = static igmp member

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show mld snooping vlan static

```
show mld snooping vlan <name> static [group | router]
```

Description

Displays static MLD snooping entries.

Syntax Description

name	Specifies a VLAN name.

Default

None.

Usage Guidelines

Use this command to display the MLD snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports will be displayed.

Example

The following command displays the MLD snooping static groups configured on VLAN *vlan101*:

```
show mld snooping vlan101 static group
```

The following is sample output for this command:

Group	Port	Flags
ff03::1:1:1	7	sa
ff03::1:1:1	15	sa

Flags: (s) Static, (a) Active

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure mld

```
unconfigure mld
```

Description

Resets all MLD settings to their default values and clears the MLD group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all MLD settings to their default values and clears the MLD group table:

```
unconfigure mld
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

This chapter describes commands used for configuring Multicast Source Discovery Protocol (MSDP). MSDP is an interdomain multicast protocol used to connect multiple multicast routing domains that run Protocol Independent Multicast-Sparse Mode (PIM-SM). MSDP speakers are configured on each PIM-SM domain. These speakers establish a peering relationship with other MSDP speakers through secured TCP connections. When the source sends traffic, the MSDP speaker learns about the source through its Rendezvous Point (RP). In turn, the RP advertises the source to its peers through Source Advertisement (SA) messages. The peers receive these advertisements and inform their RPs about the presence of the active source in the other PIM-SM domain, which triggers the normal PIM operation in the corresponding domains.

MSDP Licensing

MSDP is supported on all platforms running ExtremeXOS 12.0 (or higher) with the Core license (or higher). Refer to the section [“Software Licenses and Feature Packs” on page 46](#) for more information about licensing.

Summary of MSDP Software Features

The following MSDP software features are supported:

- MSDP mesh-group support
- Virtual Router (VR) aware
- MSDP cache server (Source-Active [SA] request/response mechanism)
- SA policing/filtering on a peer basis.
- Default MSDP peer and associated policy
- Multiple default MSDP peers
- Export of locally registered multicast sources
- Peer authentication using TCP MD5
- SNMP MIB access (read-only) on all scalar and tabular objects (as specified by RFC-4624)

For more conceptual information about configuring and managing MSDP, see the *ExtremeXOS Concepts Guide*.

clear msdp counters

```
clear msdp counters {peer <remoteaddr> | peer all | system} {vr <vrname>}
```

Description

This command resets the MSDP counters to zero.

Syntax Description

peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
system	Clears the global MSDP counters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The clear msdp counters command clears the following MSDP counters:

- Per peer counters
 - Number of SA messages received
 - Number of SA messages transmitted
 - Number of SA request messages received
 - Number of SA request messages transmitted
 - Number of SA response messages received
 - Number of SA response messages transmitted
 - Number of SA messages received without encapsulated data
 - Number of SA messages transmitted without encapsulated data
 - Number of SA messages received with encapsulated data
 - Number of SA messages transmitted with encapsulated data
 - Number of times the MSDP peer attained an “ESTABLISHED” state
 - Number of times the peer-RPF check failed
 - Number of times the TCP connection attempt failed
 - Total number of received messages
 - Total number of transmitted messages
- Global counters
 - None defined

The `clear counters` command will also clear all MSDP counters, but it clears the counters for all other applications too.

Example

The following command clears the counters for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp counters peer 192.168.45.43
```

The following command clears the all peer and global counters:

```
clear msdp counters
```

The following command clears all counters for a particular peer:

```
clear msdp counters peer 192.168.32.45
```

The following command clears the counters of all MSDP peers:

```
clear msdp counters peer all
```

The following command clears the global counters:

```
clear msdp counters system
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

clear msdp sa-cache

```
clear msdp sa-cache {{peer} <remoteaddr> | peer all} {group-address <grp-addr>} {vr <vrname>}
```

Description

This command purges all SA cache entries and notifies the PIM that the SA cache is empty.

Syntax Description

<grp-addr>	Specifies the IP address and subnet mask of the multicast group you want to clear. All SA cache entries that match the specified group address are removed from the database.
peer all	Specifies all MSDP peers. All matching SA cache entries from all peers are removed from the database.
<remoteaddr>	Specifies the IP address of the MSDP peer. All matching SA cache entries learned from the specified peer are removed from the database.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

MSDP receives SA messages periodically. So, after clearing SA cache entries from the local database, MSDP relearns those entries during the next advertisement from its peer.

Example

The following command clears SA cache records for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp sa-cache peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp default-peer

```
configure msdp [{peer} <remoteaddr> | peer all] default-peer {default-peer-policy <filter-name>} {vr <vrname>}
```

Description

This command configures a default or static RPF peer from which all MSDP SA messages are accepted. To remove the default peer, enter the `configure msdp no-default-peer` command.

Syntax Description

<filter-name>	Specifies the name of the policy filter associated with the default peer. The peer will be the default peer for all SA entries that are permitted by the policy filter. If an SA message is allowed by the policy filter, it will be accepted. Otherwise, the SA message has to go through the regular RPF-check. The static peer RPF check is the last step in peer RPF algorithm. So, if an SA message is denied by the default peer policy, ultimately the SA message will be rejected by MSDP.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no static RPF peer is configured.

The “default-peer-policy” keyword specifies the name of the policy filter associated with the default peer. You can configure multiple default peers with different policies. If no policy is specified, then the current peer is the default RPF peer for all SA messages.

Usage Guidelines

Configuring a default peer simplifies peer-RPF checking of SA messages. If the peer-RPF check fails, the default peer rule is applied to see if the SA messages should be accepted or rejected.

If a default peer policy is specified, the peer is the default peer only for the (Source, Group), or (S, G), that satisfies the policy. If the policy is not specified, then the default peer is used for all (S, G, RP).

You can configure multiple default peers on an MSDP router; however all default peers must either have a default policy or not. A mix of default peers, with a policy and without a policy, is not allowed.

When configuring multiple default peer rules, follow these guidelines:

- When you enter multiple `default-peer` commands with the `default-peer-policy` keyword, you can use all the default peers at the same time for different RP prefixes.
- When you enter multiple `default-peer` commands without the `default-peer-policy` keyword, you can use a single active peer to accept all SA messages. If that peer goes down, then the next configured default peer accepts all SA messages. This configuration is typically used at a stub site.

You can use the following policy attributes in a default peer policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 as the default peer policy for “sales”:

```
configure msdp peer 192.168.45.43 default-peer default-peer-policy sales
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp max-rejected-cache

```
configure msdp max-rejected-cache <max-cache> {vr <vrname>}
```

Description

Configures the maximum limit on rejected SA cache entries that an MSDP router will store in its database.

Syntax Description

<max-cache>	Specifies the maximum number of rejected SA cache entries that the MSDP router will store in its database. To remove the limit, enter 0 (zero) for the <max-cache> value.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the maximum cache entries stored is zero. That is, rejected SA cache entries are not stored. Any SA cache entries that are stored and not refreshed for six minutes are removed.

Usage Guidelines

SA cache are rejected because of:

- Peer-RPF failure
- Policy denied

When a previously rejected SA cache entry is accepted because of an RP reachability change or policy rule change, the rejected SA cache entry is moved to the accepted SA cache list.

By default, rejected SA cache entries are discarded. You can configure a limit for rejected cache entries to store them, which will help debug/diagnose some issues; however, it consumes extra memory.

Example

The following command sets the maximum rejected cache limit to 100 for an MSDP router:

```
configure msdp max-rejected-cache 100
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp mesh-group

```
configure msdp [{peer} <remoteaddr> | peer all] mesh-group [<mesh-group-name> | none] {vr <vrname>}
```

Description

This command configures an MSDP peer to become a member of a mesh-group. To remove a peer from a mesh-group, enter the `none` CLI keyword for the mesh-group.

Syntax Description

<mesh-group-name>	Specifies the name of the MSDP mesh-group.
none	Removes a peer from a mesh-group.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Any SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Mesh-groups achieve two goals:

- Reduce SA message flooding
- Simplify peer-RPF flooding

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 to become a member of a mesh-group called “intra”:

```
configure msdp peer 192.168.45.43 mesh-group intra
```


History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp no-default-peer

```
configure msdp [{peer} <remoteaddr> | peer all] no-default-peer {vr  
<vrname>}
```

Description

This command removes a default peer.

Syntax Description

no-default-peer	Removes a default peer.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes all MSDP peers:

```
configure msdp peer all no-default-peer
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp originator-id

```
configure msdp originator-id <ip-address> {vr <vrname>}
```

Description

Configures the originator ID for an MSDP router. The originator ID is the RP address you want to use (instead of the default) in locally originated SA messages.

Syntax Description

<ip-address>	Specifies the RP address to use in locally originated SA messages. To unconfigure an originator ID (that is, to use the default RP address), enter the IP address 0.0.0.0.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the RP address is used as the originator ID in locally originated SA messages.

Usage Guidelines

Use this command to override the default RP address used in SA messages. Because only RPs and MSDP border routers originate SAs, there are times when it is necessary to change the ID used for this purpose. The originator ID address must be one of the interface addresses on the MSDP router.

You can configure the MSDP originator ID only when MSDP is disabled globally.

To remove an originator ID, enter the IP address 0.0.0.0.

Example

The following command configures the originator ID for an MSDP router:

```
configure msdp originator-id 10.203.134.1
```

The following command unconfigures the originator ID for an MSDP router:

```
configure msdp originator-id 0.0.0.0
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp password

```
configure msdp [{peer} <remoteaddr> | peer all] password [none |
{encrypted} <tcpPassword>] {vr <vrname>}
```

Description

This command configures a TCP MD5 password for an MSDP peer. This command enables TCP MD5 authentication for a MSDP peer. When a password is configured, MSDP receives only authenticated MSDP messages from its peers. All MSDP messages that fail TCP MD5 authentication are dropped.

Syntax Description

encrypted	Encrypts the password for MD5 authentication. To improve security, the password displays in encrypted format and cannot be seen as simple text. Additionally, the password is saved in encrypted format.
none	Removes the previously configured password.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<tcpPassword>	Specifies the password to use for MD5 authentication at the TCP level. The password must be an ASCII string with a maximum of 31 characters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, TCP MD5 authentication is disabled for the MSDP peer.

Usage Guidelines

We recommend that you enable TCP MD5 authentication for all MSDP peers to protect MSDP sessions from attacks. You can execute this command only when the MSDP peer is disabled or when MSDP is globally disabled on that VR.

Example

The following command configures a password for the MSDP peer with the IP address 192.168.45.43, which automatically enables TCP MD5 authentication:

```
configure msdp peer 192.168.45.43 password test123
```

The following command removes the password:

```
configure msdp peer 192.168.45.43 password none
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp peer description

```
configure msdp peer <remoteaddr> description {<peer-description>} {vr
<vrname>}
```

Description

Configures a name or description for an MSDP peer. This text is for display purposes only.

Syntax Description

<remoteaddr>	Specifies the IP address of the MSDP peer.
<peer-description>	Specifies the name or description of the MSDP peer. The maximum is 63 characters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no name or description is specified.

Usage Guidelines

Use this command to configure a name or description to make an MSDP peer easier to identify. The description is visible in the output of the `show msdp peer` command.

To remove the description, use this command without a description string.

Example

The following command configures the name “internal_peer” to an MSDP peer:

```
configure msdp peer 192.168.45.43 description internal_peer
```

The following command removes the description from an MSDP peer:

```
configure msdp peer 192.168.45.43 description
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp sa-cache-server

```
configure msdp sa-cache-server <remoteaddr> {vr <vrname>}
```

Description

Configures the MSDP router to send SA request messages to the MSDP peer when a new member becomes active in a group.

Syntax Description

<remoteaddr>	Specifies the IP address of the MSDP peer from which the local router requests SA messages when a new member becomes active in a group, and MSDP has no cache entry for the group in the local database.
<vrname>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

You can use this command to force a new member of a group to learn the current active multicast sources in a connected PIM-SM domain that are sending to a group. The router will send SA request messages to the specified MSDP peer when a new member joins a group and MSDP doesn't have a cache entry for that group in the local database. The peer replies with the information in an SA cache response message.



NOTE

An MSDP peer must exist before it can be configured as an SA cache server. The `configure msdp sa-cache-server` command accepts the value for <remoteaddr> only if it is an existing peer's IP address.

Example

The following command configures an MSDP cache server:

```
configure msdp sa-cache-server 172.19.34.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp sa-filter

```
configure msdp [{peer} <remoteaddr> | peer all] sa-filter [in | out]
[<filter-name> | none] {vr <vrname>}
```

Description

This command configures an incoming or outgoing policy filter for SA messages.

Syntax Description

<filter-name>	Specifies the name of the policy associated with an SA filter. To remove an SA filter, enter the "none" CLI keyword for <filter-name>.
in	Associates the SA filter with inbound SA messages.
out	Associates the SA filter with outbound SA messages.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA filter is configured for an MSDP peer. That is, incoming and outgoing SA messages are not filtered.

Usage Guidelines

This command configures an SA filter such that only a specified set of SA messages are accepted or sent to a peer. Note that an SA filter does not adversely impact the flow of SA request and response messages.

To remove an SA filter, enter the "none" CLI keyword for <filter-name>.

You can use the following policy attributes in an SA filter policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command configures an incoming SA messages filter on an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-filter in allow_229
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp sa-limit

```
configure msdp [{peer} <remoteaddr> | peer all] sa-limit <max-sa> {vr
<vrname>}
```

Description

This command allows you to limit the number of SA entries from an MSDP peer that the router will allow in the SA cache. To allow an unlimited number of SA entries, use 0 (zero) as the value for <max-sa>.

Syntax Description

max-sa	Specifies the maximum number of SA entries from an MSDP peer allowed in the SA cache. To specify an unlimited number of SA entries, use 0 (zero) as the value for <max-sa>.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA entry limit is set. The router can receive an unlimited number of SA entries from an MSDP peer.

Usage Guidelines

You can use this command to prevent a distributed denial of service (DOS) attack. We recommend that you configure an MSDP SA limit on all MSDP peer sessions. Note that a rejected SA cache entry is not included in the number of SA cache entries received from a peer.

Example

The following command configures the SA entry limit of 500 for the MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-limit 500
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp source-interface

```
configure msdp [{peer} <remoteaddr> | peer all] source-interface
[<ipaddress> | any] {vr <vrname>}
```

Description

This command configures the source interface for the MSDP peer TCP connection.

Syntax Description

any	Specifies to use any interface as one end of the TCP connection. The source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection. Basically, this command removes the previously configured source interface of the MSDP peer.
<ipaddress>	Specifies the IP address of the MSDP router interface to use on one end of a TCP connection. The <ipaddress> must be one of the MSDP router interface addresses; otherwise, the command fails and an error message displays.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, the source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection.

Usage Guidelines

You must first disable MSDP or the MSDP peer before using this command. We recommend that you configure a source interface for MSDP peers that are not directly connected. We also recommend using the loopback address as the MSDP peer connection endpoint.

Example

The following command configures a source interface for an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 source-interface 60.0.0.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp timer

```
configure msdp [{peer} <remoteaddr> | peer all] timer keep-alive <keep-  
alive-sec> hold-time <hold-time-sec> {vr <vrname>}
```

Description

The command configures the keep-alive and hold timer intervals of the MSDP peers.

Syntax Description

<hold-time-sec>	Specifies the hold timer interval in seconds, in the range of 3 through 75 seconds.
<keep-alive-sec>	Specifies the keep-alive timer interval in seconds, in the range of 1 through 60 seconds.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the:

- Keep-alive timer interval is 60 seconds.
- Hold timer interval is 75 seconds.
- SA timer interval is 60 seconds.

Usage Guidelines

You can use this command only when either MSDP or the MSDP peer is disabled. The hold timer interval must be greater than the keep-alive timer interval.

Example

The following command configures the keep-alive and hold timer intervals for the MSDP peer 55.0.0.83:

```
configure msdp peer 55.0.0.83 timer keep-alive 30 hold-time 60
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure msdp ttl-threshold

```
configure msdp [{peer} <remoteaddr> | peer all] ttl-threshold <ttl> {vr
<vrname>}
```

Description

Configures the limit to which multicast data packets are sent in SA messages to an MSDP peer. If the time-to-live (TTL) in the IP header of an encapsulated data packet exceeds the TTL threshold configured, encapsulated data is not forwarded to MSDP peers.

Syntax Description

all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer on which to configure a TTL threshold.
<ttl>	Specifies the TTL value. The range is 0 through 255. To restore the default value, enter a TTL value of 0 (zero).
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

The default value is zero, meaning all multicast data packets are forwarded to the peer regardless of the TTL value in the IP header of the encapsulated data packet.

Usage Guidelines

This command allows you to configure a TTL value to limit multicast data traffic.

Example

The following command configures a TTL threshold of 5:

```
configure msdp peer 192.168.45.43 ttl-threshold 5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure pim border

```
configure pim <vlan_name> border
```

Description

Configures a PIM VLAN as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

MSDP is used to connect multiple multicast routing domains. A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN.

Example

The following command configures a PIM border on a VLAN called "vlan_border":

```
configure pim vlan_border border
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

create msdp mesh-group

```
create msdp mesh-group <mesh-group-name> {vr <vrname>}
```

Description

Creates an MSDP mesh-group.

Syntax Description

<mesh-group-name>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, there is no MSDP mesh-group.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Create a mesh-group to:

- Reduce SA message flooding
- Simplify peer-RPF flooding

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group, which reduces SA message flooding.

Example

The following command creates a mesh-group called “verizon”:

```
create msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

create msdp peer

```
create msdp peer <remoteaddr> {remote-as <remote-AS>} {vr <vrname>}
```

Description

Creates an MSDP peer.

Syntax Description

<remoteaddr>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
<remote-AS>	Specifies the autonomous system (AS) number of the MSDP peer. This number is used for display purposes and during peer-RPF checks.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The AS number for the peer is used in peer-RPF checks. The AS number provided by the user is treated as the AS number in which the peer resides. If you do not specify an AS number, BGP is used to determine the AS number for the peer. If the peer is reachable by other routes, BGP will not be able to determine the AS number for the peer.

Example

The following command creates an MSDP peer:

```
create msdp peer 192.168.45.43 remote-as 65001
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

delete msdp mesh-group

```
delete msdp mesh-group <mesh-group-name> {vr <vrname>}
```

Description

Removes an MSDP mesh-group.

Syntax Description

<mesh-group-name>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Mesh-groups are used to achieve two goals:

- Reduce SA message flooding
- Simplify peer-RPF flooding

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Use the `delete msdp mesh-group` command only if you created a mesh-group that you want to remove. By default, there is no MSDP mesh-group.

Example

The following command removes a mesh-group called “verizon”:

```
delete msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

delete msdp peer

```
delete msdp peer [all | <remoteaddr>] {vr <vrname>}
```

Description

Deletes an MSDP peer.

Syntax Description

all	Deletes all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The AS number for the peer is used in peer-RPF checks. The AS number provided by the user is treated as the AS number in which the peer resides. If you do not specify an AS number, BGP is used to determine the AS number for the peer. If the peer is reachable by other routes, BGP will not be able to determine the AS number for the peer.

Example

The following command deletes an MSDP peer:

```
delete msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable msdp

```
disable msdp {vr <vrname>}
```

Description

Disables MSDP on a virtual router.

Syntax Description

<vrname>	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to disable MSDP on a virtual router.

Example

The following command disables MSDP on a virtual router:

```
disable msdp
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable msdp data-encapsulation

```
disable msdp data-encapsulation {vr <vrname>}
```

Description

Disables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages.

Usage Guidelines

N/A.

Example

The following command disables multicast data packet encapsulation:

```
disable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable msdp export local-sa

```
disable msdp export local-sa {vr <vrname>}
```

Description

Disables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups. Use this command to disable it.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

Example

The following command disables the advertisement of local sources:

```
disable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable msdp peer

```
disable msdp [{peer} <remoteaddr> | peer all] {vr <vrname>}
```

Description

Configures the administrative state of an MSDP peer.

Syntax Description

all	Disables all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer to disable.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

Use this command to administratively disable MSDP peers to stop exchanging SA messages.

Example

The following command disables an MSDP peer:

```
disable msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable msdp process-sa-request

```
disable msdp [{peer} <remoteaddr> | peer all] process-sa-request {vr
<vrname>}
```

Description

This command configures a router to reject SA request messages from a specified peer or all peers.

Syntax Description

peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from all peers.

Usage Guidelines

Use this command to configure the router to reject SA request messages from a specified peer or all peers.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command disables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
disable msdp peer 192.168.45.43 process-sa-request
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable msdp

```
enable msdp {vr <vrname>}
```

Description

Enables MSDP on a virtual router.

Syntax Description

<vrname>	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to enable MSDP on a virtual router.

Example

The following command enables MSDP on a virtual router:

```
enable msdp
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable msdp data-encapsulation

```
enable msdp data-encapsulation {vr <vrname>}
```

Description

Enables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages.

Usage Guidelines

Enable data encapsulation to handle bursty sources.

Example

The following command enables multicast data packet encapsulation:

```
enable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable msdp export local-sa

```
enable msdp export local-sa {export-filter <filter-name>} {vr <vrname>}
```

Description

Enables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

<filter-name>	Specifies the policy to associate with the export of local sources. No policy is specified by default.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

You can use the following policy attributes in an export policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Please note that the syntax for “multicast-group”, “multicast-source,” and “pim-rp” are the same as for the “nlnri” policy attribute.

```
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any]/<mask-length>
{exact}
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any] mask <mask>
{exact}
```

An example of an MSDP policy file follows:

```
entry allow_internal_rp {
    if match any {
        multicast-group 234.67.89.0/24;
        multicast-source 23.123.45.0/24;
        pim-rp 10.203.134.5/32;
    } then {
        permit;
    }
}
entry deny_local_group239 {
    if match any {
        multicast-group 239.0.0.0/8;
        multicast-source 23.123.45.0/24;
    } then {
        deny;
    }
}
entry allow_external_rp_172 {
    if {
        multicast-group 234.172.0.0/16;
    } then {
        permit
    }
}
# deny remaining entries
```

Example

The following command enables the advertisement of local sources:

```
enable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable msdp peer

```
enable msdp [{peer} <remoteaddr> | peer all] {vr <vrname>}
```

Description

Configures the administrative state of an MSDP peer.

Syntax Description

all	Enables all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer to configure.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

You must use this command to administratively enable the MSDP peers before they can establish peering sessions and start exchanging SA messages.

Example

The following command enables an MSDP peer:

```
enable msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable msdp process-sa-request

```
enable msdp [{peer} <remoteaddr> | peer all] process-sa-request {sa-  
request-filter <filter-name> } {vr <vrname>}
```

Description

This command configures MSDP to receive and process SA request messages from a specified peer or all peers. If an SA request filter is specified, only SA request messages from those groups permitted are accepted. All others are ignored.

Syntax Description

<filter-name>	Specifies the name of the policy filter associated with SA request processing.
peer all	Specifies all MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from peers.

Usage Guidelines

Use this command to configure the router to accept all or just some SA request messages from peers. If no policy is specified, all SA request messages are accepted. If a policy is specified, only SA request messages from those groups permitted are accepted, and all others are ignored.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command enables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
enable msdp peer 192.168.45.43 process-sa-request sa-request-filter intra_domain
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show msdp

```
show msdp {vr <vrname>}
```

Description

This command displays global configuration and run-time parameters for MSDP.

Syntax Description

<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to verify the global configuration parameters of MSDP.

Example

The following command displays global configuration and run-time parameters for MSDP:

```
show msdp
```

The following is sample output from this command:

MSDP Enabled	: Yes	VR-Name	: VR-Default
Originator RP Addr	: 60.0.0.5	SA Cache ageout Int	: 360
Store SA Cache	: Yes	SA Cache Server	: 55.0.0.83
Export Local SAs	: Yes	Export SA filter	: Not configured
Max Rejected Cache	: 1200	Encapsulate data	: Yes
Num of Rejected SAs	: 0	Total Num of SAs	: 255

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show msdp memory

```
show msdp memory {detail | <memoryType>}
```

Description

This command displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process.

Syntax Description

detail	Displays detailed statistics for all memory types.
<memoryType>	Displays statistics for a particular memory type.

Default

N/A.

Usage Guidelines

Use this command to view and diagnose the memory utilization of the MSDP process.

Example

The following displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process:

```
show msdp memory
```

The following is sample output from this command:

```
MSDP Memory Information
-----
```

```
Bytes Allocated: 79792  AllocFailed: 0  OversizeAlloc: 0
Current Memory Utilization Level: GREEN
```

```
Memory Utilization Statistics
-----
```

Size	16	32	48	64	80	96	128	256	1024	4096	8192	12288
Used Blocks	0	0	256	263	3	0	2	0	1	0	0	4
peer	0	0	0	0	0	0	0	0	0	0	0	4
mesh-group	0	0	0	3	0	0	0	0	0	0	0	0
sa-node	0	0	0	255	0	0	0	0	0	0	0	0
sa-entry	0	0	255	0	0	0	0	0	0	0	0	0
vr-node	0	0	0	0	0	0	0	0	1	0	0	0
rt-cache	0	0	0	5	0	0	0	0	0	0	0	0
rp-node	0	0	1	0	0	0	0	0	0	0	0	0

client	0	0	0	0	0	0	2	0	0	0	0	0
misc	0	0	0	0	3	0	0	0	0	0	0	0

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show msdp mesh-group

```
show msdp [mesh-group {detail} | {mesh-group} <mesh-group-name>] {vr
<vrname>}
```

Description

This command displays configuration information about MSDP mesh-groups.

Syntax Description

detail	Displays detailed information about MSDP mesh-groups.
<mesh-group-name>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to display configuration information about MSDP mesh-groups, as follows:

- For summary information, enter the `show msdp mesh-group` command.
- For detailed information, enter the `show msdp mesh-group detail` command.
- For detailed information about a specific mesh-group, enter the `show msdp mesh-group <name>` command.

Example

The following command displays the peer count for a mesh-group:

```
show msdp mesh-group
```

The following is sample output from this command:

MeshGroupName	PeerCount
external	0
internal	0
msdp_mesh	4

The following command displays detailed information about a mesh-group called "msdp_mesh":

```
show msdp mesh-group "msdp_mesh"
```

The following is sample output from this command:

```
Mesh Group Name      : msdp_mesh      Num of Peers : 4
Peers                : 54.172.168.97   55.0.0.83      124.56.78.90
                     221.160.90.228
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show msdp peer

```
show msdp [peer {detail} | {peer} <remoteaddr>] {vr <vrname>}
```

Description

This command displays configuration and run-time parameters about MSDP peers.

Syntax Description

detail	Displays detailed information about MSDP peers.
<remoteaddr>	Specifies the IP address of the MSDP peer.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to verify the configuration and run-time parameters for MSDP peers, as follows:

- For summary information, enter the `show msdp peer` command.
- For detailed information for all peers, enter the `show msdp peer detail` command.
- For detailed information for a specific peer, enter the `show msdp peer <remoteaddr>` command.

Example

The following command displays configuration and run-time parameters for MSDP peers:

```
show msdp peer
```

The following is sample output from this command:

Peer Address	AS	State	Up/Down	Resets	SA_Cnt	Name
-d 54.172.168.97	14490	DISABLED	00:31:36	0	0	test
*e 55.0.0.83	100	ESTABLISHED	00:21:04	1	0	to-Hawaii
-d 124.56.78.90	2345	DISABLED	00:31:36	0	0	
-d 221.160.90.228	23456	DISABLED	00:31:36	0	0	

Flags: (*) default peer, (d) disabled, (e) enabled

The following command displays output from an MSDP peer with the IP address 55.0.0.83:

```
show msdp peer 55.0.0.83
```

The following is sample output from this command:

```

MSDP Peer      : 55.0.0.83 (to-xxxxx)
Enabled        : Yes
Source Address  : 55.0.0.5(auto)
Default Peer    : Yes (active)
Process In Request : Yes
Maximum SA Limit : 4000
Input SA Filter : not configured
State          : ESTABLISHED
Local Port     : 35600
In Total Msgs  : 24
In SA Msgs     : 0
In SA Req Msgs : 0
In SA Resp Msgs : 0
Time since Last Msg : 00:00:54
Connection Attempts : 3
RPF Fails      : 0
AS Number      : 100
TTL Threshold   : 8
Default Peer Filter : not configured
In Request filter : test
Mesh Group      : msdp_mesh
Output SA Filter : not configured
Uptime/Downtime : 00:21:35
Remote Port     : 639
Out Total Msgs  : 23
Out SA Msgs     : 22
Out SA Req Msgs : 0
Out SA Resp Msgs : 0
Hold Tmr Exp in : 00:00:21
Entered Established : 1
Output Queue Size : 0

```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show msdp sa-cache

```
show msdp [sa-cache | rejected-sa-cache] {group-address <grp-addr>}
{source-address <src-addr>} {as-number <as-num>} {originator-rp
<originator-rp-addr>} {local} {peer <remoteaddr>} {vr <vrname>}
```

Description

This command displays the SA cache database. The following quadruplet per SA cache entry displays: [Group, Source, originating RP, and peer]. In addition, information about the following displays: the cache uptime, aging, whether sources are local or remote, etc.

Syntax Description

<as-num>	Displays all SA cache that originated from the specified Autonomous System (AS) number.
<grp-addr>	Displays the SA cache within the specified group address range.
<originator-rp-addr>	Displays all SA cache entries that were originated by the specified rendezvous point.
local	Displays locally originated SA cache entries only.
<remoteaddr>	Displays the SA cache entries received from the MSDP peer with the specified IP address.
<src-addr>	Displays the SA cache within the specified source address range.
<vrname>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to view and troubleshoot the SA cache database. There are various filtering criteria you can use to display just a subset of the SA cache database. The following are some of the criteria, which you can use together or separately, to display information about the SA cache:

- Filtering on the group address range
- Filtering on the source address range
- Filtering on the originator rendezvous point address
- Filtering of the advertising MSDP peer
- Locally originated SA cache
- Rejected SA cache

Example

The following command displays the SA cache database:

```
show msdp sa-cache
```

The following is sample output from this command:

Group Address	Source Address	Originator	Peer Address	Age/Ageout In
235.100.200.1	10.20.30.1	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.2	10.20.30.2	60.0.0.5	192.0.0.16	00:44:24/05:16
235.100.200.3	10.20.30.3	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.4	10.20.30.4	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.5	10.20.30.5	60.0.0.5	55.0.0.5	00:44:24/05:01
235.100.200.6	10.20.30.6	60.0.0.5	178.54.67.23	00:44:24/05:17
235.100.200.7	10.20.30.7	60.0.0.5	112.234.213.12	00:44:24/05:43
235.100.200.8	10.20.30.8	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.9	10.20.30.9	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.10	10.20.30.10	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.11	10.20.30.11	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.12	10.20.30.12	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.13	10.20.30.13	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.14	10.20.30.14	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.15	10.20.30.15	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.16	10.20.30.16	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.17	10.20.30.17	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.18	10.20.30.18	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.19	10.20.30.19	60.0.0.5	0.0.0.0	00:44:25/00:00

Number of accepted SAs : 255

Number of rejected SAs : 0

Flags: (a) Accepted, (f) Filtered by policy, (r) RPF check failed

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

unconfigure msdp sa-cache-server

```
unconfigure msdp sa-cache-server {vr <vrname>}
```

Description

Removes the MSDP SA cache server.

Syntax Description

<vrname>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

Use this command to remove the MSDP SA cache server you specified with the `configure msdp sa-cache-server` command.

Example

The following command removes the MSDP SA cache server:

```
unconfigure msdp sa-cache-server
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

unconfigure pim border

```
unconfigure pim <vlan_name> border
```

Description

Unconfigures a PIM VLAN that has been configured as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

By default, no PIM VLANs are configured as border VLANs.

Usage Guidelines

A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN. Use the `unconfigure pim border` command to remove the border functionality of the specified PIM VLAN.

Example

The following command unconfigures a PIM border on a VLAN called “vlan_border”:

```
unconfigure pim vlan_border border
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

This appendix describes commands for:

- Downloading and using a new switch software image
- Saving, uploading, and downloading switch configuration information
- Downloading and installing a new BootROM image and switch rebooting

The switch software *image* contains the executable code that runs on the switch. An image comes preinstalled from the factory. The image can be upgraded by downloading a new version from a Trivial File Transfer Protocol (TFTP) server on the network. If you have a modular switch, you can also download a new version from the external compact flash memory card installed in the external compact flash slot of the Management Switch Fabric Module (MSM).

A switch can store up to two images; a primary and a secondary image. You can download a new image into either one of these, and you can select which image will load on the next switch reboot.

The *configuration* is the customized set of parameters that you have selected to run on the switch. As you make configuration changes, the new settings are stored in run-time memory. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

The switch can store multiple user-defined configuration files, each with its own file name. By default, the switch has two pre-named configurations: a primary and a secondary configuration. You can select to which configuration you want the changes saved, or you can save the changes to a new configuration file. You can also select which configuration will be used on the next switch reboot.

The BootROM initializes certain important switch variables during the switch boot process. In specific situations, you can upgrade the BootROM on the Summit family of switches, SummitStack, and the BlackDiamond 10808 switch by download from a TFTP server on the network. On the BlackDiamond 8800 series switches, you can upgrade the firmware, including the BootROM, when you upgrade the software image. On the BlackDiamond 12800 series switches, you can upgrade the firmware, including the BootROM, when you use the `install firmware` command.

configure firmware

```
configure firmware [auto-install | install-on-demand]
```

Description

Configures how the BlackDiamond 8800 series switch performs a system firmware upgrade.

Syntax Description

auto-install	Specifies ExtremeXOS to automatically upgrade the firmware if the software detects a newer firmware image is available. The switch does not prompt you to confirm the firmware upgrade.
install-on-demand	Specifies the switch to prompt you to upgrade the firmware when ExtremeXOS determines that a newer firmware image is available. This is the default behavior.

Default

The default is install-on-demand.

Usage Guidelines

Use the `configure firmware [auto-install | install-on-demand]` and `install firmware` commands to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in the BlackDiamond 8800 series switch.

Firmware images are bundled with ExtremeXOS software images. ExtremeXOS automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the ExtremeXOS image when you:

- Download a new version of ExtremeXOS to the active partition.
- Install a new module into an active chassis.

After a firmware image upgrade, messages are sent to the log.

If you select the `auto-install` parameter, you are not prompted to confirm the firmware upgrade. Whenever ExtremeXOS determines a newer firmware image is available, the firmware is automatically upgraded.

If you use the default configuration `install-on-demand`, you have the opportunity to cancel the firmware upgrade. If you install a new software image, and a new firmware image is available, the switch prompts you to upgrade the firmware. Enter `y` to upgrade the firmware image. Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.

If you download a new image and new firmware images are available, you see messages similar to the following:

```
download image bd8800-11.3.0.10.xos
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes

Downloading to MSM-
A.....
.....
Image will be installed to the active partition, a reboot required. Do you want to
continue? (y/n) Yes
Do you want to save configuration changes to primary.cfg? (y or n) Yes
Saving configuration on primary MSM ..... done!
Installing to primary partition!

Installing to MSM-
A.....
.....
.....
.....
.....
.....
Image installed successfully

Installing version 1.0.0.16 of the MSM bootrom(s). Do you want to continue? (y/n)
Yes

Installing version 1.0.0.24 of the IO module bootrom(s). Do you want to continue?
(y/n) Yes

Installing version 2.4 of the PSU control module firmware. Do you want to continue?
(y/n) Yes

Installing bootrom...

MSM bootrom(s) installed successfully

Installing bootrom...

IO module bootrom(s) installed successfully

Installing firmware...

PSU controller firmware installed successfully

...
```

Displaying BootROM and Firmware Versions. To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the [show version](#) command.

Recovering From a Corrupted BootROM. If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset

(R) holes on the BlackDiamond 8800 MSM and applying pressure. For more information, please refer to the *Extreme Networks Consolidated XOS Hardware Installation Guide*.

Example

The following command automatically upgrades the firmware when a newer firmware image is present without prompting you to confirm the upgrade:

```
configure firmware auto-install
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switch.

download bootrom

```
download bootrom [[<ipaddress> | <hostname>] <filename> {{vr} <vrname>}
| memorycard <filename>] {msm <slotid>}
```

On a SummitStack:

```
download bootrom [[<ipaddress> | <hostname>] <filename> {{vr} <vrname>}
| memorycard <filename>] {slot <slotid>}
```

Description

Downloads a BootROM image after the Summit family switch or the BlackDiamond 10808 switch has booted. The downloaded image replaces the BootROM in the onboard FLASH memory.

On a SummitStack, the command downloads a BootROM image to a node in the stack.

Syntax Description

ipaddress	Specifies the IP address of the TFTP server.
hostname	Specifies the hostname of the TFTP server.
memorycard	Specifies that the BootROM image should be obtained from the external compact flash memory card. NOTE: This parameter is available only on the BlackDiamond 10808 switch.
vrname	Specifies the name of the virtual router. NOTE: The Summit family of switches do not support user-created virtual routers (VRs).
filename	Specifies the name of the file that contains the BootROM image.
slotid	Specifies the MSM where the BootROM image should be downloaded. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on the BlackDiamond 10808 switch or a SummitStack. In a SummitStack, the slotid specifies the slot number of the node. On a SummitStack, the slotid specifies the node to which the image should be downloaded.

Default

N/A.

Usage Guidelines



NOTE

The Summit family of switches does not support user-created VRs.

Upgrade the BootROM image only when asked to do so by an Extreme Networks technical representative.

Prior to downloading the BootROM image on the switch, you must download the image you received from Extreme Networks to a TFTP server on your network. On the BlackDiamond 10808 switch, you can also download the image to the external compact flash memory card.

Specify the `ipaddress` or `hostname` parameters to download a BootROM image from a TFTP server on the network.

The BootROM image file is a `.xbr` file, and this file contains the executable code.

If this command does not complete successfully it could prevent the switch from booting. In the event the switch does not boot properly, some boot option functions can be accessed through a special Bootloader menu (see the *ExtremeXOS Concepts Guide*).

Use of the `hostname` parameter requires that DNS be enabled.

Displaying the BootROM Versions. To display the BootROM version for the switch and all of the modules and PSU controllers installed in a modular switch, use the `show version` command.

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)

- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

BlackDiamond 10808 Switch Only. Specify `memorycard` to download a BootROM image from the external compact flash memory card installed in the external compact flash slot of the MSM. Use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and place the image onto the card. For more information about installing the external compact flash memory card into the MSM, see the *Extreme Networks Consolidated XOS Hardware Installation Guide*.

SummitStack Only. You can issue this command only from the Master node. When provided, the file to be downloaded has to be compatible with the type of Summit in the specified slot if the user confirms that the image is to be installed.

If no slot number is provided and the user confirms installation in the case of download, an attempt is made to install the bootrom image on all active nodes. The bootrom image will not be installed on any node if the bootrom image specified is not compatible with all active nodes.

Example

The following command downloads a bootROM image from the tftp server `tftphost` from the file `bootimage`:

```
download bootrom tftphost bootimage
```

History

This command was first available in ExtremeXOS 11.0.

The `msm` parameter was added in ExtremeXOS 11.1.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on the Summit family of switches (whether or not included in a SummitStack) and the BlackDiamond 10808 switch.

download image

Using TFTP

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>}
| memorycard <filename>] {<partition>} {msm <slotid>}
```

To download an image to a stack:

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>}
| memorycard <filename>] {<partition>} {slot <slot number>}
```

Description

Downloads a new version of the ExtremeXOS software image.

The image file can be downloaded using TFTP which is not a secure method or SFTP and SCP2 which are secure methods. The procedure using TFTP begins above and for SFTP/SCP2 [on page 2369](#).

Syntax Description

hostname	Specifies the hostname of the TFTP server from which the image should be obtained.
ipaddress	Specifies the IP address of TFTP server from which the image should be obtained.
memorycard	Specifies that the image should be obtained from the external compact flash memory card. NOTE: This parameter is available only on modular switches.
filename	Specifies the filename of the new image.
vrname	Specifies the name of the virtual router. NOTE: The BlackDiamond 8800 series switch and the Summit family of switches do not support user-created VRs.
partition	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
slotid	Specifies the MSM where the software image should be downloaded. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.
slot number	Specifies the slot where the software image should be downloaded. The value may be from 1 to 8. NOTE: This parameter is available only on stackable switches in a stack.

Default

Stores the downloaded image in the partition selected by the [use image](#) command.

Using SFTP and SCP2

SFTP and SCP2 provide secure methods of downloading the ExtremeXOS software image files, *.xos or *.xmod. You can use one of three procedures:

- From the switch, running the command SCP2. connect to and “get” from a remote server. This is similar to the `download image` command.
- From outside the switch, connect to the switch which is acting as the server and “put” from the remote server. There is no TFTP equivalent for this method.
 - Using SFTP, or
 - Using SCP2.

Examples of these procedures are shown [on page 2378](#).

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Prior to downloading an image on the switch, you must download the image you received from Extreme Networks to a TFTP server on your network. On a modular switch, you can also download the image to the external compact flash memory card.



NOTE

Unlike ExtremeWare, the `download image` command in ExtremeXOS causes the switch to use the newly downloaded software image during the next switch reboot. To modify or reset the software image used during a switch reboot, use the `use image` command.

Specify the `ipaddress` or `hostname` parameters to download an image from a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

Core Software Images. The switch can store up to two core images: a primary and a secondary. When downloading a new image, you select which partition (primary or secondary) to install the new image. If you do not specify a partition, the software image is downloaded and installed into the current (active) partition. If you want to install the software image to the alternate partition, you must specify that partition before downloading the image.

Image Filenames. The software image file can be an .xos file, which contains an ExtremeXOS core image, or an .xmod file, which contains an ExtremeXOS modular software package. Modular software packages have additional functionality that supplement a core image.

You can identify the appropriate image or module for your platform based on the filename of the image. [Table 32](#) lists the filename prefixes for each platform:

Table 32: Filename prefixes

Platform	Filename Prefixes
BlackDiamond 12800 series	bd12K-
BlackDiamond 10808	bd10K-
BlackDiamond 8810	bd8800-
BlackDiamond 8806	bd8800-
Summit family	summitX-

For example, if you have a BlackDiamond 8806 switch, download image filenames with the prefix bd8800-.

Displaying the Software Image Versions. To display the software image version running on the switch, use the `show version` or `show switch` commands.

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Messages Displayed by the Switch. When you download a new image, you see the following message:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel)
```

Do one of the following:

- Enter `y` if you want to install the image after download.
- Enter `n` if you want to install the image at a later time.
- Press `[Enter]` if you want to cancel the download.

Beginning with ExtremeXOS 11.6, if you have an expired service contract and attempt to download a new image, the switch displays the following message:

```
Service contract expired, please renew it to be able to download the new software
image.
```

If you see this message, you must renew your service contract to download and install the new image. For more information about service contracts, see [Chapter 1, “Command Reference Overview.”](#)

Core Dump Messages. If you configure the switch to write core dump (debug) files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if you have a modular switch with an external memory card installed with space available, transfer the files to the external memory card. On the Summit family of switches, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

The switch displays a message similar to the following and prompts you to take action:

```
Core dumps are present in internal-memory and must be removed before this download can
continue. (Please refer to documentation for the “configure debug core-dumps” command
for additional information)
Do you want to continue with download and remove existing core dumps? (y/n)
```

Enter `y` to remove the core dump files and download the new software image. Enter `n` to cancel this action and transfer the files before downloading the image.

For information about configuring and sending core dump information, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

Modular Switches Only. Specify `memorycard` to download a an image from the external compact flash memory card installed in the external compact flash slot of the MSM. Use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer’s instructions to access the compact flash card and place the image onto the card. For more information about installing the external compact flash memory card into the MSM, see the *Extreme Networks Consolidated ExtremeXOS Hardware Installation Guide*.

SummitStack Only. You can issue this command only from the Master node.

If a slot is not specified, the image is downloaded to every node in the Active Topology. If a slot is specified, the image is downloaded to that slot only.

If all nodes to be downloaded are not running the same partition, the command is not executed and following message is displayed:

```
Error: all nodes do not have the same image partition selected.
```

If all nodes to be downloaded have the same partition selected but the EXOS is currently running from the selected partition, the command is not executed and the following message is displayed:

Error: the image partition selected must not be the active partition.

Downloading a New Image

The following assumes you have already downloaded the image to a network TFTP server or external memory card. The information in this section provides more detailed information for downloading a new image to your switch.



NOTE

Always refer to the most recent version of the ExtremeXOS Installation and Release Notes for the most current instructions.

Step 1—Verifying the Virtual Router. If you loaded the image onto an external compact flash (modular switches only), proceed to step 2.

If you loaded the image onto a TFTP server, use one of the following ping commands to confirm which virtual router reaches your TFTP server:

```
ping vr vr-Mgmt <host>
ping vr vr-Default <host>
```

At least one of these commands must successfully reach your TFTP server for you to download the image. After verifying the virtual router that reaches your TFTP server, specify that virtual router when you download the image.

Step 2—Viewing the Partition. To view your selected and booted partition, use the following command:

```
show switch
```

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition.

Step 3—Selecting the Partition. There are two ways to specify the partition. To specify the partition when downloading and installing the image, use the following command:

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>} | memorycard
<filename>] {<partition>} {msm <slotid>}
```

Or to specify the partition before downloading the image, use the following command:

```
use image {partition} <partition> {msm <slotid>}
```

Step 4—Downloading and Installing the Image. To download the image, use the appropriate, previously described, download image command.

Downloading an ExtremeXOS core image. An ExtremeXOS core image uses the file extension .xos.

- If you have an expired service contract and attempt to download a new image, you see the following message:

Service contract expired, please renew it to be able to download the new software image.

If you see this message, you must renew your service contract to proceed.

- If you do not have an expired service contract, before the download begins the switch asks if you want to install the image immediately after the download is finished.
 - If you download and install the image immediately to the active partition, the switch automatically reboots after the image is installed. This operation is not permitted on a SummitStack.
 - If you download the image to the active partition and do not immediately install the image, you do not need to reboot the switch until you use the image. This operation is not permitted on a SummitStack.
 - If you install the image to the inactive partition, you do not need to reboot the switch until you use the image. To install an image on a SummitStack, you must install to the inactive partition.

Enter `y` to install the image after download. Enter `n` to install the image at a later time.

If you download and install the software image on the active partition, the switch automatically reboots after the download and installation is completed. The following message appears when downloading and installing on the active partition (except on SummitStack):

```
Image will be installed to the active partition, a reboot required. Do you want
to continue? (y or n)
```

Enter `y` to continue the installation and reboot the switch. Enter `n` to cancel.

If you install the image at a later time, the image is still downloaded and saved to the switch, but you must use the following command to install the software and reboot the switch:

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

Where `fname` specifies the filename of the new, downloaded image.

Downloading an ExtremeXOS module image. An ExtremeXOS module image has functionality that supplements a core image. You download and install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module *bd10K-11.0.0.25-ssh.xmod* can only be installed onto the core image *bd10K-11.0.0.25.xos*.

To install a module to the inactive partition, use the `download image` command to download the module to the inactive partition, and use the `install image` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the inactive partition must be of the same version as the module. When you make the inactive partition active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition (except on SummitStack), use the `download image` command to download the module to the active partition, and use the `install image` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and command line interface (CLI) commands, of the module will be available.

Performing a Hitless Upgrade—Modular Switches Only. Hitless upgrade is a mechanism that allows you to upgrade the ExtremeXOS software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime
- Reducing the amount of traffic lost

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

You must have two MSMs installed in your switch to perform a hitless upgrade. With two MSMs installed in the switch, one assumes the role of primary and the other assumes the role of backup. The primary MSM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary MSM also synchronizes its configurations with the backup MSM which allows the backup to take over the management functions of the primary.



NOTE

If you download an image to the backup MSM, the image passes through the primary MSM before the image is downloaded to the backup MSM.

Before performing a hitless upgrade, review the following list to confirm that your system supports hitless upgrade:

- BlackDiamond 10808 switch—Both MSMs are running ExtremeXOS 11.1 or later. Earlier versions of ExtremeXOS do not support hitless upgrade.
- BlackDiamond 8800 series switch with MSM-G8X and 8800 original modules installed—Both MSMs are running ExtremeXOS 11.4 or later.
- BlackDiamond 8800 series switch with a mix of MSM-G8X, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later.
- BlackDiamond 8800 series switch with a mix of MSM-G8X, MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.
- BlackDiamond 8800 series switch with a mix of MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.
- BlackDiamond 12804 switch—Both MSMs are running ExtremeXOS 11.4 or later.

To perform a hitless upgrade, do the following:

- 1 View current switch information using the following command:

```
show switch
```

Determine your selected and booted partition, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized.

Output from this command indicates, for each MSM, the selected and booted images and if they are in the primary or the secondary partition. The selected image partition indicates which image will be used at the next reboot. The booted image partition indicates the image used at the last reboot. It is the active partition.

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

- 2 Select the partition to download the image to and download and install the new ExtremeXOS software on the backup MSM using the following command:

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>} | memorycard  
<filename>] {<partition>} {msm <slotid>}
```

**NOTE**

If the backup MSM is installed in slot B, specify msm B. If the backup MSM is installed in slot A, specify msm A.

- If you have an expired service contract and attempt to download a new image, you see the following message:

Service contract expired, please renew it to be able to download the new software image.

If you see this message, you must renew your service contract to proceed.

- If you do not have an expired service contract, before the download begins the switch asks if you want to install the image immediately after the download is finished. If you download and install the image immediately to the active partition, the switch automatically reboots the MSM.

- If you download and install the software image on the active partition, the switch automatically reboots the MSM after the image is installed. The following message appears when downloading and installing on the active partition:

Image will be installed to the active partition, a reboot required. Do you want to continue? (y or n)

Enter `y` to continue the installation and reboot the switch. Enter `n` to cancel.

- If you download and install the software image on the non-active partition, you do not need to reboot the switch.

- If you install the image at a later time, use the following command to install the software:

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

- 3 Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

- 4 Initiate failover from the primary MSM to the backup MSM using the following command:

```
run msm-failover
```

When you failover from the primary MSM to the backup MSM, the backup becomes the new primary, runs the newly downloaded software, and provides all of the switch management functions.

If you have a BlackDiamond 8800 series switch and the new ExtremeXOS image supports hitless upgrade but is not compatible with the current running I/O module image (the I/O version numbers do not match), you cannot perform a hitless upgrade.

The switch displays a warning message similar to the following:

```
WARNING: The other MSM operates with a different version of I/O module image.
If you continue with the MSM failover, all I/O modules will be reset.
```

Are you sure you want to failover? (y/n)

You can either continue the upgrade or cancel the action. If you continue the upgrade, the primary MSM downloads the new image to the I/O module and reboots.

- 5 Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

- 6 Select the partition to download the image to and download and install the new ExtremeXOS software on the new backup MSM (this was the original primary MSM) using the following command:

```
download image [<hostname> | <ipaddress>] <filename> {vr <vrname>} msm <slotid>
```



NOTE

If the new backup MSM is installed in slot A, specify msm A. If the new backup MSM is installed in slot B, specify msm B.

Before the download begins, the switch asks if you want to install the image immediately after the download is finished. If you download and install the image immediately to the active partition, the switch automatically reboots the MSM.

- If you download and install the software image on the active partition, the switch automatically reboots the MSM after the image is installed. The following message appears when downloading and installing on the active partition:

```
Image will be installed to the active partition, a reboot required. Do you
want to continue? (y or n)
```

Enter **y** to continue the installation and reboot the switch. Enter **n** to cancel.

- If you download and install the software image on the non-active partition, you do not need to reboot the switch.
 - If you install the image at a later time, use the following command to install the software:
- 7 Verify that the new backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

- 8 Optionally, initiate failover from the new primary MSM to the new backup MSM using the following command:

```
run msm-failover
```

When you failover from the new primary MSM to the new backup MSM, this optional step restores the switch to the original primary and backup MSM.

- 9 Optionally, confirm that the failover is successful by checking the current state of the MSMs using the following command:

```
show switch
```

You can also perform a hitless upgrade on ExtremeXOS modular software packages (.xmod files). To perform a hitless upgrade of a software package, you must install the core software image first, and the version number of the modular software package must match the version number of the core image that it will be running with. For more information about hitless upgrade, see the *ExtremeXOS Concepts Guide*.

Hitless Upgrade Caveats for the BlackDiamond 8800 Series Switch Only. The following is a summary of hitless upgrade caveats for only the BlackDiamond 8800 series switch:

- If you are running ExtremeXOS 11.5 or earlier, **do not** attempt to install builds prior to 11.6.1 on an MSM-48, and do not use the [synchronize](#) command if either the primary or secondary slot has a pre-11.6.1 image and the secondary MSM is an MSM-48. On a BlackDiamond 8800 series switch with MSM-48 modules installed, these versions of software are incompatible and cannot exist on MSMs installed in the same chassis even during the hitless upgrade process.



CAUTION

If a pre-11.6 image is loaded onto an MSM-48, the MSM will no longer boot correctly and it may require a system recovery process. If this occurs, please contact the Extreme Networks Technical Support Assistance Center (TAC) for further instructions on recovering your system.

- If you attempt a hitless upgrade between major releases, the switch warns you that the upgrade is not hitless. You can either continue the upgrade or cancel the action. If you continue the upgrade, the primary MSM downloads the new image to the I/O module and reboots them.
- If you are running ExtremeXOS 11.4 or earlier, **do not** attempt to perform a hitless upgrade to ExtremeXOS 11.5 or later. On the BlackDiamond 8800 series switch with MSM-G8X modules installed, these versions of software are incompatible and cannot exist on MSMs installed in the same chassis even during the hitless upgrade process. If attempted, the backup MSM enters the non-operational state.

To recover from the non-operational state, do the following:

- a From the primary MSM, use the [synchronize](#) command to return the MSMs to the same version of software.
- b To confirm the MSMs are synchronized, use the [show switch](#) command.

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

After you recover from the non-operational state and confirm the MSMs are synchronized, perform a normal code upgrade to install and upgrade the image on the switch. For more information, see the sections starting with [“Step 1—Verifying the Virtual Router” on page 2372](#) through [“Downloading an ExtremeXOS core image” on page 2372](#).



NOTE

Extreme XOS 11.6 introduces support for the BlackDiamond 8800 series MSM-48 module. If your switch is running an earlier version of ExtremeXOS, you must upgrade to ExtremeXOS 11.6 prior to installing the MSM-48 in your chassis.

ExtremeXOS 11.5 introduces support for the BlackDiamond 8800 a-series and e-series modules. If your switch is running ExtremeXOS 11.4 or earlier, you must upgrade to ExtremeXOS 11.5 to operate the modules.

Hitless upgrade is not supported between major releases. Do not attempt to perform a hitless upgrade. To upgrade the switch from ExtremeXOS 11.4 or earlier to ExtremeXOS 11.5 or later, reboot the switch after downloading and installing the new image to both installed MSMs.

Examples

Using TFTP

The following command downloads the switch software image from the TFTP server at 10.10.15.04, from the file named *bd10K-11.0.0.25.xos*:

```
download image 10.10.15.04 bd10K-11.0.0.25.xos
```

Although the example command specifies a BlackDiamond 10808 image, the output is similar on all modular switches. On a modular switch, if you download an image into the active partition, you see output similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes
```

```
Downloading to
MSM-A.....
Downloading to MSM-B.....
Image will be installed to the active partition, a reboot required. Do
you want to continue? (y/n) Yes
Installing to primary partition!
```

```
Installing to
MSM-
B.....
.....
.....
.....
.....
Installing to
MSM-
A.....
.....
.....
.....
```

If you answer yes to installing the image, the switch reboots upon completion of the installation.

The following command downloads the switch software image from the TFTP server at 10.10.15.04, from the file named *summitX450-11.2.0.15.xos*:

```
download image 10.10.15.04 summitX450-11.2.0.15.xos
```

If you download an image into the active partition, you see output similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes
```

```
Downloading to
Switch.....
Image will be installed to the active partition, a reboot required. Do
you want to continue? (y/n) Yes
Installing to primary partition!
.....
.....
.....
.....
.....
```

If you answer yes to installing the image, the switch reboots upon completion of the installation.

Using SFTP and SCP2

The following commands show the three procedures

- From the Switch using SCP2

```
(debug) SummitX450-24x.5 # scp2 root@10.203.133.13:/tmp/summitX-12.0.0.4-ssh.xmod
summitX-12.0.0.4-ssh.xmod
Download to /scratch/summitX-12.0.0.4-ssh.xmod on switch
Connecting to 10.203.133.13...
root@10.203.133.13's password:
Fetching /tmp/summitX-12.0.0.4-ssh.xmod to /scratch/summitX-12.0.0.4-ssh.xmod
/tmp/summitX-12.0.0.4-ssh.xmod 100% 997KB 996.9KB/s 00:00
Please wait validating image summitX-12.0.0.4-ssh.xmod ,this will take
approximately 30 seconds ...
* (debug) SummitX450-24x.6 #
```

- From Outside the Switch using SFTP

```
sftp admin@10.203.3.231
Connecting to 10.203.3.231...
admin@10.203.3.231's password:
sftp> put summitX-12.0.0.4-ssh.xmod Examplesftp.xmod
Uploading summitX-12.0.0.4-ssh.xmod to /config/Examplesftp.xmod
sftp>
```

- From Outside the Switch using SCP2

```
scp2 summitX-12.0.0.4-ssh.xmod admin@10.203.3.231:Examplescp.xmod
Keyboard-interactive:
Keyboard-interactive authentication
Enter password for admin:
summitX-12.0.0.4-ssh.xmod | 997kB | 997kB/s | TOC: 00:00:01 | 100%
[~/cvs/branch-fixes_v1200b4.c4/exos/cougar/release]$
```

Following download, the output from a `show log` command is similar to the following for the last two procedures—the two from outside the switch.

```
show log
04/03/2007 14:32:29.31 <Info:AAA.LogSsh> Got Image file Example.xmod
04/03/2007 14:32:29.31 <Info:AAA.LogSsh> Validating Image file, this could take
approximately 30 seconds.. Example.xmod
04/03/2007 14:32:30.89 <Info:AAA.LogSsh> Image file Example.xmod successfully
validated
can now install image
```

History

This command was first available in ExtremeXOS 10.1.

The `memorycard` option was added in ExtremeXOS 11.0.

The `msm` parameter was added in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

install bootrom

```
install bootrom <fname> {reboot} {msm <slotid>}
```

On a SummitStack use:

```
install bootrom <fname> {reboot} {slot <slotid>}
```

Description

Installs a new version of the ExtremeXOS BootROM image on a Summit family switch, SummitStack, or BlackDiamond 10808 switch.

Syntax Description

fname	Specifies the BootROM image file.
reboot	Reboots the switch after the image is installed.
slotid	Specifies the MSM where the BootROM image should be installed. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on the BlackDiamond 10808 switch and SummitStack. On a SummitStack, the slotid specifies the node on which the BootROM image should be installed.

Default

N/A.

Usage Guidelines

When you download a BootROM image, the system asks if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The BootROM image file is an .xbr file, and this file contains the executable code.

Displaying the BootROM Versions. To display the BootROM version for the switch and all of the modules and PSU controllers installed in a modular switch, use the `show version` command.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)

- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

SummitStack Only. You can issue this command only from the Master node.

Example

The following command installs the bootrom image file bd10K-1.0.1.5-bootrom.xbr:

```
install bootrom bd10K-1.0.1.5-bootrom.xbr
```

On the BlackDiamond 10808 switch, you see output similar to the following:

```
Installing bootrom...
Writing MSM B bootrom
.....
.....
.....
Verifying Flash contents...
.....
.....
.....
.....
bootrom written.
Writing MSM A bootrom
.....
.....
.....
Verifying Flash contents...
.....
.....
.....
.....
bootrom written.

Bootrom installed successfully
```

On the Summit series switch, you see output similar to the following:

```
Installing bootrom...
Writing bootrom
.....
.....
.....
Verifying Flash contents...
.....
.....
.....
.....
```

```
bootrom written.
```

```
Bootrom installed successfully
```

History

This command was first available in ExtremeXOS 11.0.

The `msm` parameter was added in ExtremeXOS 11.1.

From ExtremeXOS 12.0, this command is supported on a stack. The `slot` parameter is added. The `slot` parameter is applicable only when the switch is in a stack.

Platform Availability

This command is available only on the Summit family of switches, SummitStack, and the BlackDiamond 10808 switch.

install firmware

```
install firmware
```

Description

Installs the firmware bundled with the ExtremeXOS image on the BlackDiamond 8800 series switches and the BlackDiamond 12800 series switches.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

On the BlackDiamond 8800 series switch, use the `install firmware` command to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in the switch. On the BlackDiamond 12800 series switches, use the `install firmware` command to upgrade the BootROM images on the MSM and I/O modules and the firmware on the power fan controllers (PFCs) installed in the switch.

Firmware images are bundled with ExtremeXOS software images.

On the BlackDiamond 8800 series switch, ExtremeXOS automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the ExtremeXOS image. You can also use the `install firmware` command to compare the firmware images. On the BlackDiamond 12800 series switches, you must use the `install firmware` command for ExtremeXOS to compare the firmware images.

Before using the `install firmware` command, wait until the `show slot` command indicates the MSMs and I/O modules are operational. When the modules are operational, use the `install firmware` command.

BlackDiamond 8800 Series Switch. The switch scans the I/O and MSM modules and the PSU controllers for a possible firmware upgrade. If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade.

- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `<cr>` to cancel the upgrade. After a firmware image upgrade, messages are sent to the log.

The PSU controller firmware is used immediately after it is installed without rebooting the switch. The new BootROM and firmware overwrite the older versions flashed into the hardware. Use the `reboot` command to reboot the switch and activate the new BootROM and firmware.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the firmware may be corrupted and need to be recovered. ExtremeXOS automatically recovers corrupted firmware; however, the time it takes for the switch to boot-up may increase.

The switch displays status messages after you use the `install firmware` command. The output varies depending upon your platform and the software version running on your system.

**NOTE**

If the information in the most current version of the ExtremeXOS Installation and Release Notes differs from the information in this section, follow the release notes.

Sample Output—BlackDiamond 8800 Series Switch. The following is sample output from the BlackDiamond 8800 series switch:

```
Installing version 1.0.1.7 of the MSM bootrom(s). Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Do you want to save configuration changes to primary.cfg? (y or n) Yes
Saving configuration on primary MSM ..... done!
Installing version 1.0.1.7 of the IO module bootrom(s). Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 2.11 of the PSU control module firmware. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes

Installing bootrom...

MSM bootrom(s) installed successfully
MSM bootrom(s) will be activated upon next MSM reboot

Installing bootrom...

IO module bootrom(s) installed successfully
IO module bootrom(s) will be activated upon next IO module reboot

Installing firmware...
PSU controller firmware installed successfully
```

Additional Behavior—BlackDiamond 8800 Series Switch Only. During a firmware upgrade, the switch prompts you to save your configuration changes to the current, active configuration. Enter `y` to save your configuration changes to the current, active configuration. Enter `n` if you do not want to save your changes.

In earlier versions of ExtremeXOS, you are required to immediately reboot the system after a firmware upgrade. In ExtremeXOS 11.3.3 and later, the system displays a message informing you that the new firmware image will be activated the next time you reboot the system.

Use the `configure firmware [auto-install | install-on-demand]` command to configure how the switch performs a system firmware upgrade. If you select the `auto-install` parameter, you are not prompted to confirm the firmware upgrade. If you use the default configuration `install-on-demand`, you can cancel the firmware upgrade.

Power over Ethernet (PoE) firmware is always automatically upgraded or downgraded to match the operational code image. This configuration is not applicable to PoE firmware.

Recovering From a Corrupted BootROM—BlackDiamond 8800 Series Switch Only. If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset (R) holes on the BlackDiamond 8800 MSM and applying pressure. For more information, please refer to the *Extreme Networks Consolidated XOS Hardware Installation Guide*.

BlackDiamond 12800 Series Switches.



NOTE

When you use the `install firmware` command to install the new BootROM and firmware images, all of the I/O modules enter the Down state. Upgrade the firmware when you can schedule a downtime for your network.

If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade:

- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `<cr>` to cancel the upgrade.

The PFC firmware is used immediately after it is installed without rebooting the switch. The new BootROM and firmware overwrite the older versions flashed into the hardware. If you upgrade the firmware image, a reboot is required. Use the `reboot` command to reboot the switch.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the firmware may be corrupted and need to be recovered. ExtremeXOS automatically attempts to recover corrupted firmware; however, in some situations user intervention is required.

After a firmware image upgrade, messages are sent to the log.

Sample Output—BlackDiamond 12804 Switch. The output was modified in ExtremeXOS 11.6. The following is sample output from the BlackDiamond 12804 switch:

```
Would You Like to Install MSM and IOB Firmware?
(y - yes, n - no, <cr> - cancel) Yes
No default configuration database has been selected to boot up the system.
Save configuration will set the new configuration as the default database.

Do you want to save configuration changes to primary.cfg? (y or n) No
Would You Like to Install PFC Firmware?
(y - yes, n - no, <cr> - cancel) Yes
Installing firmware...

PSU controller firmware installed successfully
EXOS may take up to 10 seconds to recognise that a new PSU firmware image has been
installed
Installing firmware...
No Firmware Install Required For Slot 1
No Firmware Install Required For Slot 2
Installing Firmware to Slot 5...
No Firmware Install Required For Slot 6

IO module firmware installed successfully
```

```
IO module firmware will be activated upon next IO module reboot
Installing firmware...
Installing Firmware to MSM...
```

```
MSM firmware installed successfully
MSM firmware will be activated upon next MSM reboot
```

Additional Behavior—BlackDiamond 12800 Series Switches Only. To install the new BootROM and firmware, wait until the `show slot` command indicates the MSM(s) and I/O modules are operational. When the modules are operational, use the `install firmware` command to install the BootROM and/or firmware only on modules that have older BootROM or firmware images installed.

**NOTE**

Upgrading the firmware on a BlackDiamond 12800 series switch can take a minimum of 6 minutes per slot.

Displaying BootROM and Firmware Versions. To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the `show version` command.

Example

The following command installs the newer firmware image(s):

```
install firmware
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond 8800 series switches and the BlackDiamond 12800 series switches.

install image

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

On a SummitStack use:

```
install image <fname> {<partition>} {slot <slot number>} {reboot}
```

Description

Installs a new version of the ExtremeXOS software image.

Syntax Description

fname	Specifies the software image file.
partition	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
slotid	Specifies the MSM where the software image file should be installed. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slotid specifies the node to which the BootROM image should be installed.
reboot	Reboots the switch after the image is installed.

Default

N/A.

Usage Guidelines

When you download a software image, you are asked if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The software image file can be an .xos file, which contains an ExtremeXOS core image, or an .xmod file, which contains additional functionality to supplement a core image.

SummitStack Only . You can issue this command only from a Master node. The slot parameter is available only on a stack.

Displaying the Software Image Versions. To display the software image version running on the switch, use the `show version` or `show switch` commands.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local, remember the requirements listed above.

Installing an ExtremeXOS core image. If you install the software image on the active partition (This is not supported on SummitStack), you must reboot the switch. A message similar to the following appears when installing the image on the active partition:

```
Image will be installed to the active partition, a reboot required. Do you want
to continue ? (y/n)
```

Enter `y` to continue the installation and reboot the switch. Enter `n` to cancel.

If you install on the inactive partition, you can continue to run the currently booted image, but to run the newly installed image, you will need to set the boot partition with the `use image {partition} <partition> {msm <slotid>}` command and reboot the switch.

Installing an ExtremeXOS module image. An ExtremeXOS module image has functionality that supplements a core image. You will install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module `bd10K-11.0.0.25-ssh.xmod` can only be installed onto the core image `bd10K-11.0.0.25.xos`.

To install a module to the inactive partition, use the `install image` command to install the module. Remember, the core image on the inactive partition must be of the same version as the module. When you make the inactive partition active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition, use the `install image` command to install the module. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and CLI commands, of the module will be available.

Performing a Hitless Upgrade—Modular Switches Only. If you specify the `msm` parameter on a BlackDiamond 10808, BlackDiamond 12804, or BlackDiamond 8800 series switch, you can initiate hitless upgrade between the primary and backup MSMs installed in the switch.

Hitless upgrade is a mechanism that allows you to upgrade the ExtremeXOS software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime
- Reducing the amount of traffic lost

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

Regardless of how you upgrade the software, you must:

- Review the following list to confirm that your system supports hitless upgrade:
 - BlackDiamond 10808 switch—Both MSMs are running ExtremeXOS 11.1 or later. Earlier versions of ExtremeXOS do not support hitless upgrade.
 - BlackDiamond 8800 series switch with MSM-G8X and 8800 original modules installed—Both MSMs are running ExtremeXOS 11.4 or later.
 - BlackDiamond 8800 series switch with a mix of MSM-G8X, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later.
 - BlackDiamond 8800 series switch with a mix of MSM-G8X, MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.
 - BlackDiamond 8800 series switch with a mix of MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.
 - BlackDiamond 12804 switch—Both MSMs are running ExtremeXOS 11.4 or later.
- View the current switch information to determine your selected and booted image partitions, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized using the `show switch` command
- Select the partition to use when downloading the image using the `download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>} | memorycard <filename>] {<partition>} {msm <slotid>}` or the `use image {partition} <partition> {msm <slotid>}` commands

When performing a hitless upgrade, you must first download the software to the backup MSM. If you download the image to the inactive partition, use the `install image <fname> {<partition>} {msm <slotid>} {reboot}` command to install the software image at a later time. After the software is downloaded and installed on the switch, use the `run msm-failover` command to initiate failover from the primary MSM to the backup MSM. The original primary MSM becomes the new backup MSM.

After failover is complete, download the software to the new backup MSM. Again, if you download the image to the inactive partition, use the `install image <fname> {<partition>} {msm <slotid>} {reboot}` command to install the software image at a later time.

For more detailed information about hitless upgrade, see the `download image` command.

Example

The following command installs the software image file `bd12K-12.0.1.5.xos` on a BlackDiamond 12800 series switch:

```
install image bd12K-12.0.1.5.xos
```

The following command installs the software image file `bd10K-11.1.0.20.xos` on a BlackDiamond 10808 switch:

```
install image bd10K-11.1.0.20.xos
```

The following command installs the software image file `bd8800-11.3.0.10.xos` on a BlackDiamond 8810 switch:

```
install image bd8800-11.3.0.10.xos
```

The following command installs the software image file summitX450-11.5.1.2.xos on a Summit family switch:

```
install image summitX450-11.5.1.2.xos
```

History

This command was first available in ExtremeXOS 10.1.

The `msm` parameter was added in ExtremeXOS 11.1.

The `slot` parameter was added to support SummitStack in ExtremeXOS 12.0

Platform Availability

This command is available on all platforms.

load script

```
load script <script-file>
```

Description

Loads (plays back) the ASCII-formatted configuration file on the switch.

Syntax Description

script-file	Specifies the user-defined name of the ASCII-formatted configuration file. The script file is known as the XOS CLI script file.

Default

N/A.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch and the Summit family of switches do not support user-created VRs.

After downloading the configuration file, you must load (playback) the new configuration on the switch. Use this command to load and restore the ASCII-formatted configuration file that you downloaded from the TFTP server to the switch.

An ASCII-formatted configuration file uses the .xsf file extension, not the .cfg file extension. The .xsf file extension (known as the XOS CLI script file) saves the XML-based configuration in an ASCII format readable by a text editor.

After issuing this command, the ASCII configuration quickly scrolls across the screen. The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

For more detailed information about the ASCII configuration file, including the steps involved to upload, download, and save the configuration, see the [upload configuration \[<hostname> | <ipaddress>\] <filename> {vr <vr-name>}](#) command.

Example

The following command loads the ASCII-formatted configuration named configbackup.xsf:

```
load script configbackup.xsf
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

run update

```
run update
```

Description

Activates a newly installed modular software package.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After you install a modular software package to the active partition, use this command to make the update active. This command causes the ExtremeXOS system to start the newly installed processes contained in the package, without rebooting the switch.

If you installed the package to the inactive partition, you need to reboot the switch to activate the package.

Example

The following command activates any newly installed modular software packages installed on the active partition:

```
run update
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

save configuration

```
save configuration {primary | secondary | <existing-config> | <new-config>}
```

Description

Saves the current configuration from the switch's runtime memory to non-volatile memory.

Syntax Description

primary	Specifies the primary saved configuration.
secondary	Specifies the secondary saved configuration.
existing-config	Specifies an existing user-defined configuration.
new-config	Specifies a new user-defined configuration.

Default

Saves the current configuration to the location used on the last reboot.

Usage Guidelines

The configuration takes effect on the next reboot.

Each file name must be unique and can be up to 32 characters long but cannot include any spaces, commas, or special characters.

Configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension. Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using ASCII-formatted configuration files see the [upload configuration \[<hostname> | <ipaddress>\] <filename> {vr <vr-name>}](#) and the [load script <script-file>](#) commands.

This command also displays in alphabetical order a list of available configurations. The following is sample output that displays the primary, secondary, and user-created and defined configurations ("test" and "XOS1" are the names of the user-created and defined configurations):

```
exsh.9 # save configuration
<cr>          Execute the command
primary       Primary configuration file
secondary     Secondary configuration file
<existing-config> Existing configuration file name
    "test"    "XOS1"
<new-config>  New configuration file name
```

The switch prompts you to save your configuration changes. Enter **y** to save the changes or **n** to cancel the process.

If you enter **n**, the switch displays a message similar to the following:

```
Save configuration cancelled.
```

If you enter `y`, the switch saves the configuration and displays a series of messages. The following sections provide information about the messages displayed when you save a configuration on your switch.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Saving a New Configuration. If you create and save a configuration with a new file name, the switch saves the new configuration and then prompts you to select the newly created configuration as the switch's default configuration.

The following sample output is similar to the message displayed:

```
Do you want to save configuration to test1.cfg? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test1.cfg successfully.
```

The switch then prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n)
```

Enter `y` to use the new configuration as the default configuration. Enter `n` to cancel the operation and keep using the current default, active configuration.

Saving an Existing Configuration. If you make and save changes to an existing configuration, the switch prompts you to save and override the existing configuration.

The following sample output is similar to the message displayed:

```
The configuration file test.cfg already exists.
Do you want to save configuration to test.cfg and overwrite it? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test.cfg successfully.
```

The following sample output on a SummitStack is similar to the output displayed:

```
The configuration file primary.cfg already exists.
Do you want to save configuration to primary.cfg and overwrite it? (y/N) Yes
Saving configuration on primary ..... done!
Synchronizing configuration to backup .... done!
```

```
Saving config on Standbys (Slots: 1).  
...  
Configuration saved on Standby (Slot 1): done!
```

If you override an existing configuration that is not the current default, active configuration, the switch prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system  
(primary.cfg) is different than the one just saved (test.cfg).  
Do you want to make test.cfg the default database? (y/n) No  
Default configuration database selection cancelled.
```

Enter `y` to use the updated configuration as the default configuration. Enter `n` to cancel the operation and keep using the current default, active configuration.

Example

The following command saves the current switch configuration to the configuration file named `XOS1`:

```
save configuration XOS1
```

The following command save the current switch configuration to the secondary configuration file:

```
save configuration secondary
```

History

This command was first available in ExtremeXOS 10.1.

The status messages displayed by the switch were updated in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show configuration

```
show configuration {<module-name>}
```

Description

Displays a configuration to the terminal.

Syntax Description

module-name	Specifies the name of configuration module. The term configuration module refers to feature in ExtremeXOS. By displaying a module, you can view the commands used to configure that feature. For example, to display all of the configurations that you made for only STP, specify the stp as the module-name.

Default

N/A.

Usage Guidelines

If the output scrolls off the top of the screen, you can use the `enable clipaging` command to pause the display when the output fills the screen. The default for clipaging is enabled.

Extreme Networks recommends using the `show configuration` command to view on the CLI your currently running switch configuration. These files have the .cfg file extension. Do not use a text editor to view or modify your XML-based switch configuration files.

To save the configuration file as an ASCII-formatted file, and to view it with a text editor, see the `upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}` and the `load script <script-file>` commands.

If you specify `show configuration` without any keywords, the switch displays configuration information for each of the switch modules.

You can display only the configuration of a module of interest by using the `module-name` keyword. For example, some of the modules are AAA, ACL, BGP, EDP, FDB, SNMP, and VLAN. Use [TAB]-completion to see a list.

You must have administrator access to view the output of the `show configuration` command.

Depending on the software version running on your switch, the configurations on your switch, and the type of switch you have, additional or different configuration information may be displayed.

Example

This command shows the current configuration active in the switch:

```
show configuration
```

The following is sample output from this command:

```
# Module devmgr configuration.
#
configure snmp sysName "Engineering"
configure snmp sysContact "support@extremenetworks.com, +1 888 257 3000"

#
# Module vlan configuration.
#
create virtual-router "VR-Default"
create vlan "Default"
configure vlan Default tag 1
configure vlan Default qosprofile QP1

#
# Module fdb configuration.
#
configure fdb agingtime 300
configure iparp vr VR-Control max_entries 4096
configure iparp vr VR-Control max_pending_entries 256
configure iparp vr VR-Control max_proxy_entries 256
configure iparp vr VR-Control timeout 20
enable iparp vr VR-Control checking
enable iparp vr VR-Control refresh
...
```

This command shows the current configuration of the OSPF module in the switch:

```
show configuration ospf
```

The following is sample output from this command:

```
# Module ospf configuration.
#
configure ospf routerid automatic
configure ospf spf-hold-time 3
configure ospf metric-table 10M 10 100M 5 1G 4 10G 2
configure ospf lsa-batch-interval 30
configure ospf import-policy none
configure ospf ase-limit 0
disable ospf originate-default
disable ospf use-ip-router-alert
disable ospf
disable ospf export direct
disable ospf export static
disable ospf export rip
disable ospf export e-bgp
disable ospf export i-bgp
configure ospf area 0.0.0.0 external-filter none
configure ospf area 0.0.0.0 interarea-filter none
configure ospf area 0.0.0.0 normal
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show memorycard

```
show memorycard
```

Description

Displays whether an external memory card is present in the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to verify if there is an external memory card in the switch.

Example

This command shows whether a memory card is present on the switch:

```
show memorycard
```

If you do not have an external memory card installed, the output is similar to the following:

```
Memorycard is not present.
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available only on modular switches.

synchronize

synchronize

On a SummitStack, use the following command:

```
synchronize {slot <slotid>}
```

Description

The `synchronize` command replicates all saved images and configurations from the primary MSM or node to the backup MSM or target node on a switch or SummitStack.

Syntax Description

slotid	On a SummitStack, the slotid specifies the target node that has to be synchronized with the Master node. If the slotid is omitted, the target is the backup node. NOTE: This parameter is available only on SummitStack.

Default

N/A.

Usage Guidelines

This command:

- 1 Reboots the backup MSM or target node to prepare it for synchronizing with the primary MSM or node
- 2 Performs a binary copy of the primary MSM or node to the backup MSM or target node, including the primary and secondary software images, all configurations and policies, and temporary files
- 3 Reboots the backup MSM or target node after replication is complete

During a synchronization, half of the switch fabric is lost. When the primary MSM or node finishes replicating its configurations and images to the backup MSM or target node, the full switch fabric is restored.

To use the `synchronize` command make sure your switch or SummitStack is running the following software:

- BlackDiamond 10808 switch—Both MSMs are running ExtremeXOS 11.0 or later.
- BlackDiamond 8810 switch with MSM-G8X and 8800 original modules installed—Both MSMs are running ExtremeXOS 11.1 or later.
- BlackDiamond 8806 switch with MSM-G8X and 8800 original modules installed—Both MSMs are running ExtremeXOS 11.3 or later.
- BlackDiamond 8800 series switch with a mix of MSM-G8X, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later.

- BlackDiamond 8800 series switch with a mix of MSM-G8X, MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.
- BlackDiamond 8800 series switch with a mix of MSM-48, 8800 original, and 8800 a- and e-series modules installed—Both MSMs are running ExtremeXOS 11.6 or later.

**CAUTION**

BlackDiamond 8800 Series Only—Do not attempt to install builds prior to version 11.6.1 on an MSM-48, and do not use the `synchronize` command if either the primary or secondary slot has a pre-11.6.1 image and the secondary MSM is an MSM-48. If a pre-11.6 image is loaded onto an MSM-48, the MSM will no longer boot correctly and it may require a system recovery process. If this occurs, please contact the Extreme Networks Technical Support Assistance Center (TAC) for further instructions on recovering your system.

- BlackDiamond 12804 switch—Both MSMs are running ExtremeXOS 11.4 or later.
- SummitStack—All nodes are running ExtremeXOS 12.0 or later.

When you install a backup MSM, you are not prompted to synchronize the images and the configurations from the primary. If not synchronized, the backup uses its image and the primary's configuration. This image/configuration mismatch will likely cause the switch to operate differently after failover. Use the `synchronize` command to replicate all saved images and configurations from the primary to the backup.

If you have not saved your runtime configuration, you are prompted to save it when you use the `synchronize` command. A message similar to the following appears:

```
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter `y` to save the configuration and continue with synchronizing the MSMs. Enter `n` to cancel the operation. If you enter `y`, messages similar to the following appear:

```
Saving configuration on primary MSM ..... done!
Synchronizing configuration to backup MSM .. done!
```

After the configuration has been saved and replicated to the backup MSM, synchronization begins.

After the initial reboot, if the backup MSM is not available or does not respond within 120 seconds, the synchronize operation fails.

Use the `show switch {detail}` command to verify that the backup MSM is in sync with the primary MSM.

BlackDiamond 8800 Series Switch Only. On the BlackDiamond 8800 series switch, the I/O ports on the backup MSM go down when you synchronize the MSMs. When the primary MSM finishes replicating its configurations and images to the backup MSM, the I/O ports on the backup MSM come back up.

SummitStack Only. While using the command `synchronize {slot <slot-number>}` if the slot number is provided, that slot is the target of the synchronize operation. If the slot number is not provided, the Backup node is synchronized. This command can only be executed on the Master node.

Example

The following example assumes you have already saved your runtime configuration.

The following command replicates all saved images and configurations from the master MSM to the backup MSM:

```
synchronize
```

After you enter `synchronize`, status messages similar to the following appear:

```
Synchronize will reboot the backup MSM, then overwrite all code images
and configs with a copy from the master MSM.
```

```
Synchronize currently requires ExtremeXOS version 11 or greater on
the backup MSM
```

```
DO NOT interrupt synchronize, the backup MSM may become unbootable!
```

```
OK to continue? (y/n) Yes
```

```
Rebooting Backup MSM...
```

```
NOTE: The command line is locked during synchronize
```

```
synchronizing...
```

```
synchronizing nvram...
```

```
synchronizing nvram...
```

```
synchronizing nvram...
```

```
synchronizing nvram...
```

```
synchronizing nvram...
```

```
synchronizing nvram...
```

```
synchronizing XOS...
```

```
[=====] 100% XOS
```

```
Synchronize complete - rebooting backup MSM...
```

```
BD-10808.4 #
```

History

This command was first available in ExtremeXOS 11.0.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0

Platform Availability

This command is available only on modular switches and SummitStack.

synchronize stacking

```
synchronize stacking {node-address <node-address> | slot <slotid>}
```

Description

The `synchronize stacking` command replicates all saved images and configurations from the primary node to the backup node on a SummitStack.

Syntax Description

slotid	On a SummitStack, the slotid specifies the target node that has to be synchronized with the Master node. If the slotid is omitted, the target is the backup node. NOTE: This parameter is available only on SummitStack.
node address	The MAC address of the node to be synchronized. This must be a valid node in the active topology.

Default

N/A.

Usage Guidelines

This command:

- 1 Reboots the target node to prepare it for synchronizing with the primary node
- 2 Performs a binary copy of the primary node to the target node, including the primary and secondary software images, all configurations and policies, and temporary files
- 3 Reboots the backup target node after replication is complete

During a synchronization, half of the switch fabric is lost. When the primary node finishes replicating its configurations and images to the target node, the full switch fabric is restored.

To use the `synchronize stacking` command make sure your switch is running the following software:

- SummitStack—All nodes are running ExtremeXOS 12.0 or later.

When you install a backup node, you are not prompted to synchronize the images and the configurations from the primary. If not synchronized, the backup uses its image and the primary's configuration. This image/configuration mismatch will likely cause the SummitStack to operate differently after failover. Use the `synchronize stacking` command to replicate all saved images and configurations from the primary to the backup.

If you have not saved your runtime configuration, you are prompted to save it when you use the `synchronize stacking` command. A message similar to the following appears:

```
Do you want to save configuration changes to primary.cfg? (y or n)
```


Enter `y` to save the configuration and continue with synchronizing the nodes. Enter `n` to cancel the operation.

After the configuration has been saved and replicated to the target node, synchronization begins.

After the initial reboot, if the target node is not available or does not respond within 120 seconds, the synchronize operation fails.

Use the `show switch {detail}` command to verify that the target node is in sync with the primary MSM.

While using the command `synchronize {slot <slot-number>}` if the slot number is provided, that slot is the target of the synchronize operation. If the slot number is not provided, the Backup node is synchronized. This command can only be executed on the Master node.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack.

unconfigure switch

```
unconfigure switch {all}
```

Description

Returns the switch configuration to its factory default settings and reboots the switch.

Syntax Description

all	Specifies that the entire configuration should be changed to the default values, including the management IP address, failsafe account, and SummitStack-specific parameters, and the switch rebooted.

Default

N/A.

Usage Guidelines

Use `synchronize stacking` to reset the configuration to factory defaults, but without erasing the configuration. This preserves users account information, date and time settings, SummitStack configuration, and so on.

Include the parameter `all` to clear the entire current configuration, including all switch and SummitStack parameters, and reboot using the last used image and factory default configuration.

The command `unconfigure switch {all}` does not clear licensing information. The license cannot be disabled once it is enabled on the switch.

For SummitStack only. The "all" option clears or resets to default values all of the following stacking configuration parameters:

- Stacking enabled
- Slot number
- Stack MAC address
- Stack node priority
- Master-capable setting
- Alternate management IP address

Example

The following command preserves the entire current configuration (but does not reload the current configuration after the switch reboots) and reboots the switch or SummitStack using the last specified saved image and factory default configuration:

```
unconfigure switch all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

uninstall image

```
uninstall image <fname> <partition> {msm <slotid>} {reboot}
```

On a SummitStack, use:

```
uninstall image <fname> <partition> {slot <slot number>} {reboot}
```

Description

Uninstalls an ExtremeXOS software package.

Syntax Description

fname	Specifies the software package to uninstall.
partition	Specifies which partition the package was installed to: primary or secondary. Select primary to remove it from the primary partition and secondary to remove it from the secondary partition.
slotid	Specifies the MSM where the package should be uninstalled. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slotid specifies the node on which the BootROM image should be installed.
reboot	Reboots the switch after the package is uninstalled.

Default

N/A.

Usage Guidelines

Use this command to uninstall a software package previously installed on the switch.

When you uninstall a software package, the switch prompts you to save your changes to your currently active configuration file:

```
Uninstallation of the EXOS module
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter **y** to save the changes to your configuration file. Enter **n** to not save the changes to your configuration file.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements previously described.

SummitStack Only. You can issue this command only from a Master node.

Example

The following command uninstalls the software package *bd10K-11.0.0.25.xos* from the secondary partition:

```
uninstall image bd10K-11.0.0.25.xos secondary
```

History

This command was first available in ExtremeXOS 11.0.

The `msm` parameter was added in ExtremeXOS 11.1.

The `slot` parameter was added to support SummitStack in ExtremeXOS 12.0

Platform Availability

This command is available on all platforms.

upload configuration

```
upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}
```

Description

Uploads the current configuration in ASCII format to a TFTP server on your network.

Syntax Description

hostname	Specifies the hostname of the TFTP server where you want to download the configuration file. You must have DNS enabled
ipaddress	Specifies the IP address of the TFTP server where you want to download the configuration file.
filename	Specifies a user-defined name for the configuration file. You must use the .xsf file extension when naming an ASCII-formatted configuration file.
vr-name	Specifies the name of the virtual router. By default the switch uses VR-Mgmt for this command. NOTE: The BlackDiamond 8800 series switch and the Summit family of switches do not support user-created VRs.

Default

Uploads the current configuration in ASCII format immediately to a TFTP server.

Usage Guidelines



NOTE

The BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches do not support user-created VRs.

Specify the `ipaddress` or `hostname` parameters to upload the current, active configuration file from the switch to a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

The uploaded ASCII file retains the CLI format. This allows you to do the following:

- Modify the configuration using a text editor, and later download a copy of the file to the same switch or to one or more different switches.
- Send a copy of the configuration file to Extreme Networks Technical Support for problem-solving purposes.

This command is not applicable to XML-based configurations. Those files use the .cfg file extension.

If you want to view your configuration in ASCII format, use the .xsf file extension (known as the XOS CLI script file) when you save the configuration file on the switch. This saves the XML-based configuration in an ASCII format readable by a text editor.

If you successfully upload the active configuration to the network TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

If the switch displays a timeout error message similar to the following:

```
failed!
Error: timeout
```

Make sure you entered the correct host name or IP address of the TFTP server

If the switch displays an unreachable network error similar to the following:

```
failed!
Error: Network is unreachable
```

Make sure you entered the correct virtual router. By default the switch uses VR-Mgmt for this command.

Summary of Steps. The following summary only describes the CLI involved to transfer the configuration and load it on the switch; it is assumed that you know how to modify the configuration file with a text editor. As previously described, to use these commands, use the .xsf file extension. These steps are not applicable to configurations that use the .cfg file extension.

To work with an ASCII-formatted configuration file, complete the following tasks:

- 1 Upload the configuration to a network TFTP server using the following command:

```
upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}
```

After the configuration file is on the TFTP server, use a text editor to the desired changes.

- 2 Download the configuration from the TFTP server to the switch using one of the following commands:

```
tftp [<host-name> | <ip-address>] -g -r <remote-file>
```

```
tftp get [<host-name> | <ip-address>] <remote-file>
```

- 3 Verify the configuration file is on the switch using the following command:

```
ls
```

- 4 Load and restore the new configuration file on the switch using the following command:

```
load script <script-file>
```

- 5 Save the configuration to the configuration database so the switch can reapply the configuration after switch reboot using the following command:

```
save configuration {primary | secondary | <existing-config> | <new-config>}
```

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

The following describes the steps in more detail.

Uploading the ASCII Configuration File To a TFTP Server. To upload the current switch configuration as an ASCII-based file to the TFTP server, use the `upload configuration` command and save the configuration with the .xsf file extension.

For example, to transfer the current switch configuration as an ASCII-based file named `meg_upload_config1.xsf` to the TFTP server with an IP address of 10.10.10.10, do the following:

```
upload configuration 10.10.10.10 meg_upload_config1.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

Downloading the ASCII Configuration File to the Switch. To download the configuration from the TFTP server to the switch, use the `tftp` command. For example, to retrieve the configuration file named `meg_upload_config1.xsf` from a TFTP server with an IP address of 10.10.10.10, you can use one of the following commands:

```
tftp 10.10.10.10 -g -r meg_upload_config1.xsf
tftp get 10.10.10.10 meg_upload_config1.xsf
```

If you successfully download the configuration to the switch, the switch displays a message similar to the following:

```
Downloading meg_upload_config1.xsf to switch... done!
```

Verifying that the ASCII Configuration File is on the Switch. To confirm that the ASCII configuration file is on the switch, use the `ls` command. The file with an `.xsf` extension is the ASCII configuration.

The following sample output contains an ASCII configuration file:

```
-rw-r--r-- 1 root 0 98362 Nov 2 13:53 Nov022005.cfg
-rw-r--r-- 1 root 0 117136 Dec 12 12:56 epicenter.cfg
-rw-r--r-- 1 root 0 68 Oct 26 11:17 mcastgroup.pol
-rw-r--r-- 1 root 0 21203 Dec 13 15:40 meg_upload_config1.xsf
-rw-r--r-- 1 root 0 119521 Dec 6 14:35 primary.cfg
-rw-r--r-- 1 root 0 96931 Nov 11 11:01 primary_11_11_05.cfg
-rw-r--r-- 1 root 0 92692 Jul 19 16:42 secondary.cfg
```

Loading the ASCII Configuration File. After downloading the configuration file, you must load the new configuration on the switch. To load and restore the ASCII configuration file, use the `load script <script-file>` command. After issuing this command, the ASCII configuration quickly scrolls across the screen.

The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

Instead of entering each command individually, the script runs and loads the CLI on the switch.

Saving the Configuration. After you load the configuration, save it to the configuration database for use by the switch. This allows the switch to reapply the configuration after a switch reboot. To save the

configuration, use the `save configuration {primary | secondary | <existing-config> | <new-config>}` command.

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

You can use any name for the configuration. For example, after loading the file meg_upload_config1.xsf, you need to save it to the switch. To save the configuration as configuration1.cfg, do the following:

```
save configuration configuration1
```

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions. This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a remote file, remember the requirements previously described.

Example

The following command uploads the current switch configuration as an ASCII-based file named configbackup.xsf to the TFTP server with an IP address of 10.10.10.10:

```
upload configuration 10.10.10.10 configbackup.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading configbackup.xsf to 10.10.10.10 ... done!
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

use configuration

```
use configuration [primary | secondary | <file_name>]
```

Description

Configures the switch to use a previously saved configuration on the next reboot.

Syntax Description

primary	Specifies the configuration file named <i>primary.cfg</i> .
secondary	Specifies the configuration file named <i>secondary.cfg</i> .
file_name	Specifies an existing user-defined configuration file name (displays a list of available user-defined configuration files).

Default

N/A.

Usage Guidelines

XML-based configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension.

Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using and saving ASCII-formatted configuration files see the `upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}` and the `load script <script-file>` commands.

There is no special significance to the primary and secondary configurations. They are just conveniences to specify the files *primary.cfg* and *secondary.cfg*.

When you configure the switch to use a previously saved configuration, the switch displays the following message:

```
The selected configuration will take effect after the next switch reboot.
```

You can create a new configuration file by saving your current switch configurations and using that file on the next reboot. For example, to create a new configuration named *test1* based on your current CLI session and switch configurations, use the following command:

```
save configuration test1
```

Tracking and Displaying Switch Configuration Files. To keep track of your configuration file names, use the `ls` command to display the files saved on your switch. Files with the .cfg extension are configuration files. In addition, you can see a list of available configuration files when you use the `use configuration` command.

The following is sample output from this command ("*test*" and "*XOS1*" are the names of the user-created and defined configurations):

```
exsh.1 # use configuration
primary      Primary configuration file
secondary    Secondary configuration file
<file-name>  Configuration file name
"test"      "XOS1"
```

You can also use the `ls` command to display a list of the current configuration and policy files in the system.

Displaying the Active Configuration. To view the currently active, running configuration, use the `show switch` command.

Local Filename Character Restrictions. This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Example

The following command specifies that the next reboot should use the saved configuration file named *XOS1.cfg*:

```
use configuration XOS1
```

The following command specifies that the next reboot should use the configuration saved in the primary partition:

```
use configuration primary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

use image

```
use image {partition} <partition> {msm <slotid>}
```

On a SummitStack, use:

```
use image {partition} <partition> {slot <slotid>}
```

Description

Configures the switch to use a saved image on the next reboot.

Syntax Description

partition	Specifies which image to use on the next reboot, the one stored on the primary partition, or the one stored on the secondary partition.
slotid	Specifies the MSM where the package should be uninstalled. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slotid specifies the node on which the BootROM image is selected.

Default

The currently booted image.

Usage Guidelines

This command specifies which image to use on the next reboot. Two images can be stored, one on the primary partition, one on the secondary partition. To view your current (active) partition and the selected partition for the next reboot or installation, use the following command:

```
show switch
```

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition. Primary indicates the saved image in the primary partition; secondary indicates the saved image in the secondary partition.

This command also specifies which partition to install images to and which partition to download images to if not specified in the `install image` and `download image` commands.

SummitStack Only. You can issue this command only from a Master node. The image to use is stored in NVRAM on all target nodes.

If a slot number is not provided, the partition is selected on all nodes in the Active Topology.

Example

Using TFTP

The following command configures the switch to use the image stored in the primary partition on the next reboot:

```
use image partition primary
```

A message similar to the following is displayed:

```
To take effect of partition change please reboot the switch!
```

History

This command was first available in ExtremeXOS 10.1.

The `msm` parameter was added in ExtremeXOS 11.1.

The `slot` parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

B Troubleshooting Commands

If you encounter problems when using your switch, ExtremeXOS provides troubleshooting commands. Use these commands only under the guidance of Extreme Networks technical personnel.

This appendix describes commands for troubleshooting your switch, including:

- Running diagnostics and displaying diagnostic test results
- Enabling and disabling the standalone Extreme Loop Recovery Protocol (ELRP) client
- Enabling and disabling debug mode for Event Management System (EMS) components

You can contact Extreme Networks Technical Support at (800) 998-2408 or (408) 579-2826.

Extreme Loop Recovery Protocol

ExtremeXOS 11.1 introduced support for the Extreme Loop Recovery Protocol (ELRP). ELRP is a feature of ExtremeXOS that allows you to prevent, detect, and recover from Layer 2 loops in the network. You can use ELRP with other protocols such as the Extreme Standby Router Protocol (ESRP) and the Ethernet Automatic Protection Switching Protocol (EAPS).

A switch running ELRP transmits multicast packets with special MAC destination address out of some or all of the ports belonging to a Virtual LAN (VLAN). All the other switches in the network treat this packet as a regular, multicast packet and flood it to all the ports belonging to the VLAN. If the packets transmitted by a switch are received back by that switch, this indicates a loop in the Layer 2 network.

Standalone ELRP gives you the ability to send ELRP packets, either periodically or on an ad hoc “one-shot” basis on a specified subset of VLAN ports. If any of these transmitted packets is received back then standalone ELRP can perform a configured action such as sending a log message to the system log file or sending a trap to the SNMP manager.

Standalone ELRP allows you to:

- Configure ELRP packet transmission on specified VLANs.
- Specify some or all the ports of VLAN for packet transmission.



NOTE

Reception of packets is not limited to any specific ports of the VLAN and cannot be configured.

- Configure transmission of ELRP packets on specified ports of a VLAN periodically with the added ability to configure the interval between consecutive timings.
- Save and restore standalone ELRP configuration across reboots.

- Request periodic or non-periodic transmission of ELRP packets on specified ports of a VLAN.

For **non-periodic** ELRP requests:

- You can specify the number of times ELRP packets must be transmitted and the interval between consecutive transmissions.
- A message is printed to the console and logged into the system log file indicating detection of network loop when ELRP packets are received back or no packets are received within the specified duration.
- There is no need to trap to the SNMP manager.

For **periodic** ELRP requests:

- If ELRP packets are received back, a message is printed to the system log file and a trap is sent to the SNMP manager indicating detection of a network loop.

Details of using ELRP with ESRP are discussed in the *ExtremeXOS User Guide* chapter, "[Extreme Standby Router Protocol](#)," and the commands used to configure ELRP with ESRP are described in [Chapter 19, "ESRP Commands."](#)

Event Management System

The Event Management System (EMS) provides enhanced features to filter and capture information generated on a switch. Details of using EMS are discussed in the *ExtremeXOS User Guide*, in the chapter, "Status Monitoring and Statistics", and the commands used for EMS are detailed in this document in [Chapter 5, "Commands for Status Monitoring and Statistics."](#)

Included in this chapter, as well as in [Chapter 5](#), are the EMS commands to enable and disable debug mode for EMS components.

clear elrp counters

```
clear elrp counters
```

Description

Clears and resets the ELRP counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics before you delete the ELRP counters. Use the [show log counters](#) command to display event statistics.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period that you are monitoring.

Example

The following command clears all switch statistics related to ELRP:

```
clear elrp counters
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure debug core-dumps

```
configure debug core-dumps [internal-memory | memorycard | off]
```

Description

Enables or disables sending debug information to the specified memory card.

Syntax Description

internal-memory	Specifies that saving debug information to the internal memory card is enabled. Use this parameter only under the guidance of Extreme Networks Technical Support personnel.
memorycard	Specifies that saving debug information to the external memory card is enabled. Use this parameter only under the guidance of Extreme Networks Technical Support personnel. NOTE: This parameter is available only on modular switches.
off	Specifies that saving debug information to the memory card is disabled. This is the default behavior.

Default

By default, the switch does not automatically send debug information to the memory card. The setting is set to `off`.

Usage Guidelines



NOTE

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

The switch sends debug information to the preinstalled internal memory card. On modular switches, you can also send debug information to the external memory card installed in the external compact flash slot of the MSM.

The switch only generates core dump files and writes them to the specified memory card in the following situations:

- If an ExtremeXOS process fails.
- When forced under the guidance of Extreme Networks Technical Support.

If you configure the switch to write core dump files to the internal memory and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, move or delete the core dump files from the internal memory. For example, if you have a modular switch with an external memory card installed with space available, transfer the files to the external memory card. On the Summit family of switches, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Modular Switches Only. Before you can enable and save debug information to the external memory card, you must install an external compact flash memory card into the external compact flash slot of the MSM. For more information about installing an external compact flash memory card, please refer to the *Extreme Networks Consolidated XOS Hardware Installation Guide*.

After you use the `eject memorycard` command and manually remove the card from the external compact flash slot of the MSM, this setting is automatically changed to `off`.

Stackables in Stack Mode. This command works only from the master node. If you enable it on stack master, it is applicable for all nodes.

Example

The following example enables a modular switch to save debug information to the external memory card:

```
configure debug core-dumps memorycard
```

The following example enables the switch to save debug information to the internal memory card:

```
configure debug core-dumps internal-memory
```

History

This command was first available in ExtremeXOS 11.1.

The `internal-memory` parameter was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure elrp-client one-shot

```
configure elrp-client one-shot <vlan_name> ports [<ports> | all] interval
<sec> retry <count> [log | print | print-and-log]
```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.

Syntax Description

vlan_name	Specifies a VLAN name.
ports	Specifies the set of VLAN ports for packet transmission.
all	Specifies all ports of this VLAN for packet transmission.
sec	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 600 seconds. The default is 1 second.
count	Specifies the number of times ELRP packets must be transmitted. The range is 1 to 255 times. The default is 3 times.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print	Specifies that a message should be printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print-and-log	Specifies that a message should be logged in the system log file and printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.

Default

- **sec**—The interval between consecutive packet transmissions is 1 second.
- **count**—The number of times ELRP packets must be transmitted is 3.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client can perform a configured action such as logging a message in the system log file or printing a log message to the console. There is no need to send a trap to the SNMP manager for non-periodic requests.



NOTE

This command is backward compatible with Extreme Networks switches running ExtremeWare. If your network contains switches running only ExtremeXOS, you can also use the `run elrp <vlan_name> {ports <ports>} {interval <sec>} {retry <count>}` to perform one-time ELRP packet transmission.

Use the `configure elrp-client periodic` command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following example starts one-time, non-periodic ELRP packet transmission on all ports of the VLAN *sales*, uses the default interval and transmission times, and sends messages to the console:

```
configure elrp-client one-shot sales ports all interval 1 retry 3 print
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure elrp-client periodic

```
configure elrp-client periodic <vlan_name> ports [<ports> | all] interval
<sec> [log | log-and-trap | trap]
```

Description

Starts periodic ELRP packet transmission on the specified ports of the VLAN using the specified interval.

Syntax Description

vlan_name	Specifies a VLAN name.
ports	Specifies the set of VLAN ports for packet transmission.
all	Specifies all ports of this VLAN for packet transmission.
sec	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 600 seconds. The default is 1 second.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
log-and-trap	Specifies that a message should be logged in the system log file and a trap message should be sent to the SNMP manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
trap	Specifies that a trap message should be sent to the SNMP manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.

Default

The interval between consecutive packet transmissions is 1 second.

Usage Guidelines

This command starts periodic ELRP packet transmission on the specified ports of the VLAN using the specified interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client can perform a configured action such as logging a message in the system log file and/or sending a trap to the SNMP manager.

Use either the `configure elrp-client one-shot` or the `run elrp` command to configure non-periodic, one-time transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following example starts periodic ELRP packet transmission on slot 3, port 2 of VLAN *marketing*, sends packet transmissions every 2 seconds, and sends messages to the log:

```
configure elrp-client periodic marketing ports 3:2 interval 2 log
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure forwarding hash-algorithm

```
configure forwarding hash-algorithm [crc16 | crc32]
```

Description

Modifies hardware table utilization by configuring the hash algorithm.

Syntax Description

crc16	Specifies the CRC16 hash algorithm. This is the default setting.
crc32	Specifies the CRC32 hash algorithm.

Default

In ExtremeXOS 11.5, the default hash algorithm is crc32

In ExtremeXOS 11.4 and earlier, the default hash algorithm is crc16.

Usage Guidelines



NOTE

Modify the hardware table hash algorithm only with the guidance of Extreme Networks technical personnel.

The switch uses a hash algorithm to decide where to store the addresses in the hardware table. The standard, default hash algorithm works well for most systems; however, for some addresses with certain patterns, the hardware may attempt to store address information in the same section of the hardware.

If you are running ExtremeXOS 11.4 or earlier and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
<Info:HAL.Ipv4Adj.Info> : adj 136.159.188.109: IP add error is Table full for new or
newly resolved ARP, egress valid
<Info:HAL.Ipv4Adj.Info> : adj 136.159.188.109: returned -17 for L3 table bucket 181
<Warn:HAL.Ipv4Mc.Warning> : Could not allocate a hardware S,G,V entry
(889f4648,effffffa,70) - hardware table resource exceeded (rv=-17).
```

If you are running ExtremeXOS 11.5 or later and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
HAL.Ipv4Adj.L3TblFull> MSM-A: IPv4 unicast entry not added. Hardware L3 Table full.
Card.Ipv4Adj.Warning> Slot 4: IPv4 unicast entry not added. Hardware L3 Table full.
HAL.Ipv4Mc.GrpTblFullEnt> MSM-A: IPv4 multicast entry (10.0.0.1,224.1.1.1,vlan 1) not
added. Hardware Group Table full.
Card.Ipv4Mc.Warning> Slot-4: IPv4 multicast entry not added. Hardware L3 Table full.
```


In the previously described situations, you can configure a different hash algorithm to select a different section of the hardware to store addresses. You must save your configuration and reboot the switch to modify the hash algorithm used by the hardware table.

Upgrading to ExtremeXOS 11.5. When you upgrade to ExtremeXOS 11.5, the hash algorithm automatically becomes crc32. For example, if you saved a configuration using an image from ExtremeXOS 11.4 or earlier with the hash algorithm set to crc16, when ExtremeXOS 11.5 loads, the hash algorithm becomes crc32. To change the hash algorithm to crc16, use the `configure forwarding hash-algorithm crc16` and save your switch configuration.

Example

The following command modifies the hardware table hash algorithm to crc16:

```
configure forwarding hash-algorithm crc16
```

The switch displays the following message to prompt you to save your configuration and reboot the switch:

```
Warning: This command will only take effect after a save and reboot
```

History

This command was first available in ExtremeXOS 11.3.2.

The default hash algorithm was changed to crc32 in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches.

disable elrp-client

```
disable elrp-client
```

Description

Disables the ELRP client (standalone ELRP) globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the ELRP globally so that none of the ELRP VLAN configurations take effect.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

Example

The following command globally disables the ELRP client:

```
disable elrp-client
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable log debug-mode

```
disable log debug-mode
```

Description

Disables debug mode. The switch stops generating debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of `debug-summary`, `debug-verbose`, or `debug-data` when configuring filters
- Target format options `process-name`, `process-id`, `source-function`, and `source-line`

Example

The following command disables debug mode:

```
disable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

eject memorycard

```
eject memorycard
```

Description

Ensures that the external memory card can be safely and manually removed from the external compact flash slot on the MSM.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After the switch writes to the external memory card, and before you can view the contents on the card, you must ensure it is safe to remove the card from the external compact flash slot on the MSM. Use the `eject memorycard` command to prepare the card for removal. After you issue the `eject memorycard` command, you can manually remove the card from the external compact flash slot on the MSM and read the data on the card.

If the `configure debug coredumps memorycard` command is in effect when you issue the `eject memorycard` command, the behavior is similar to issuing the `configure debug coredumps off` command.

For more information about removing the external memory card, please refer to the *Extreme Networks Consolidated XOS Hardware Installation Guide*.

To access and read the data on the card, use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and read the data.

Example

The following command prepares the external memory card to be removed from the external compact flash slot on the MSM:

```
eject memorycard
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on modular switches.

enable elrp-client

```
enable elrp-client
```

Description

Enables the Extreme Loop Recovery Protocol (ELRP) client (standalone ELRP) globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The ELRP client must be enabled globally in order for it to work on any VLANs.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following command globally enables the ELRP client:

```
enable elrp-client
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable log debug-mode

```
enable log debug-mode
```

Description

Enables debug mode. The switch generates debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of `debug-summary`, `debug-verbose`, or `debug-data` when configuring filters
- Target format options `process-name`, `process-id`, `source-function`, and `source-line`

Example

The following command enables debug mode:

```
enable log debug-mode
```

When you enable debug mode, the following message appears:

```
WARNING: Debug mode should only be enabled when advised by technical support,  
or when advanced diagnosis is required. Performance degradation is possible.  
Debug mode now enabled.
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

nslookup

```
nslookup <host-name-or-ip>
```

Description

Displays the IP address or hostname of the requested host.

Syntax Description

host-name-or-ip	Specifies the hostname or IP address.

Default

N/A.

Usage Guidelines

For nslookup to work, you must configure the DNS client, and the switch must be able to reach the DNS server.

Host Name and Remote IP Address Character Restrictions. This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Example

The following command looks up the IP address of a computer with the name of *bigserver.xyz_inc.com*:

```
nslookup bigserver.xyz_inc.com
```

History

This command was first available in ExtremeXOS 10.1.

Support for using an IP address to obtain the name of the host was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

run diagnostics

```
run diagnostics [extended | normal | stack-port] {slot [<slot> | A | B]}
```

Description

Runs normal or extended diagnostics on the switch, slot, node, or management module. On the Summit family of switches, this command also runs diagnostics on the stacking ports.

This command is not supported in stacking mode. But if you issue the show diagnostics command from the master node, it will show the diagnostic results for all the nodes.

Syntax Description

extended	Runs an extended diagnostic routine. Takes the ports offline, and performs extensive ASIC and packet loopback tests on all of the ports. BlackDiamond 8800 series switch only—If you have a Power over Ethernet (PoE) module installed, the switch also performs an extended PoE test, which tests the functionality of the inline power adapter.
normal	Runs a normal diagnostic routine. Takes the ports offline, and performs a simple ASIC and packet loopback test on all of the ports.
stack-port	Runs the diagnostic routine on the stack ports. NOTE: This parameter is available only on Summit family of switches that are not operating in stacking mode.
slot	Specifies the slot number of an I/O module. NOTE: This parameter is available only on modular switches.
A B	Specifies which MSM to run diagnostics on. • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Depending on your platform, use this command to run diagnostics on the switch, slot, management module, or stack port.

Running Diagnostics—BlackDiamond 10808 Switch and BlackDiamond 8800 Series Switch Only. If you run the diagnostic routine on an I/O module, that module is taken offline while the diagnostic test is performed. The ports on the module do not forward traffic. Once the diagnostic test is completed, the I/O module is automatically reset and becomes operational again.

After the switch runs the diagnostic routine, test results are saved in the module's EEPROM and messages are logged to the syslog.

**NOTE**

On the BlackDiamond 10808 switch, run diagnostics when the switch can be brought offline. The tests conducted are extensive and affect traffic that must be processed by the system CPU. The CPU processes the diagnostics whether you run diagnostics on an I/O module or a management module.

On an I/O module, the extended diagnostic routine can require significantly more time to complete, depending on the number of ports on the module.

On a management module, the module is taken offline while the diagnostics test is performed. Once the diagnostic test is completed, the MSM reboots, and becomes operational again.

**NOTE**

BlackDiamond 8810 switch—If you run diagnostics on slots 5 and 6 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM.

BlackDiamond 8806 switch—if you run diagnostics on slots 3 and 4 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM.

BlackDiamond 8800 series switch—To run diagnostics on the management portion of the master MSM, specify slot A or B.

Running Diagnostics—BlackDiamond 12800 Series Switches Only. If you run the diagnostic routine on a BlackDiamond 12800 series switch, the diagnostic routine occurs on all of the installed I/O modules and both the primary and, if present, the secondary MSM. If a backup MSM is present, the switch does not test the backup MSM.

Running diagnostics affects system operation. During a diagnostic routine, the system can be offline from 5 to 25 minutes depending on the number of modules installed and the type of diagnostic test performed.

Once the diagnostic routine is completed, the primary MSM reboots the backup MSM, if installed, and then reboots itself. The system becomes operational again.

After the switch runs the diagnostic routine, test results are saved in the each module's EEPROM and messages are logged to the syslog.

Running Diagnostics—Summit Family of Switches Only. If you run the diagnostic routine on the switch, it reboots and then performs the diagnostic test. During the test, traffic to and from the ports on the switch is temporarily unavailable. When the diagnostic test is complete, the switch reboots and becomes operational again.

To run the diagnostic routine on the stack ports, you need a dedicated stacking cable that connects stack port 1 to stack port 2, which are located at the rear of the switch. The stacking cable is available from Extreme Networks. The switch performs a hardware test to confirm that the stack ports are operational; traffic to and from the ports on the switch is temporarily unavailable. This Bit Error Rate Test (BERT) provides an analysis of the number of bits transmitted in error.

After the switch runs the diagnostic routine, test results are saved to the switch's EEPROM and messages are logged to the syslog.

To run diagnostics on a Summit switch that is in a SummitStack, first disable stacking on that switch, then restart the switch. Once restarted, log into the switch via its console port, and run diagnostics. The switch will perform the diagnostic tests and then restart. Once restarted, log into the switch via its console port and enable stacking, then reboot the switch. Once restarted, the switch will rejoin the stack.

Viewing Diagnostics. To view results of the last diagnostics test run, use the following command:

```
show diagnostics {slot [<slot> | A | B]}
```



NOTE

The slot, A, and B parameters are available only on modular switches.

If the results indicate that the diagnostic failed on a module, replace the module with another module of the same type.

If the results indicate that the diagnostic failed on the switch, contact Extreme Networks Technical Support.

Example

BlackDiamond 10808 switch and BlackDiamond 8800 series switch example

The following command runs normal diagnostics on the I/O module installed in slot 2 of either the BlackDiamond 10808 switch or the BlackDiamond 8800 series switch:

```
run diagnostics normal slot 2
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter *y* to continue and run the diagnostics. Enter *n* to cancel the operation.

BlackDiamond 12800 series switch example

The following command runs normal diagnostics on all of the installed I/O modules and the primary MSM of the BlackDiamond 12800 series switch:

```
run diagnostics normal
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter *y* to continue and run the diagnostics. Enter *n* to cancel the operation.

Summit family switch example

The following command runs normal diagnostics on the Summit family switch:

```
run diagnostics normal
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter *y* to continue and run the diagnostics. Enter *n* to cancel the operation.

The following command runs diagnostics on the stack ports on the Summit X450 family switch:

```
run diagnostics stack-port
```

If you issue this command with a console connection, the switch displays the following information. You also have the opportunity to continue or cancel the test:

```
Summit Diagnostics Mode Enabled, Starting Diagnostics....
```

```
Found X450a-24T in Motherboard
Motherboard CPLD Revision: 2
Starting stacking port diagnostics
```

```
*****
*                                                                 *
* Please connect a cable between Stack Port 1 and Stack Port 2. *
*                                                                 *
*           Press S to skip test, ENTER key to continue.         *
*                                                                 *
*****
```

Press [Enter] to continue and run the diagnostics. Enter *s* to cancel the operation.

If you continue with diagnostics, the switch displays messages similar to the following:

```
Stack Port 1 and Stack Port 2 BERT
.....
.....
```

```
Stacking ports
```

```
Port 1 (Device 0 - Device port 26)
```

```
Lane 0 PASSED.
```

```
Lane 1 PASSED.
```

```
Lane 2 PASSED.
```

```
Lane 3 PASSED.
```

```
Port 2 (Device 0 - Device port 27)
```

```
Lane 0 PASSED.
```

```
Lane 1 PASSED.
```

```
Lane 2 PASSED.
```

```
Lane 3 PASSED.
```

```
DIAGNOSTIC PASS: run test bert stacking
```

```
Summit Diagnostics completed, rebooting system...
```

If you issue this command with a Telnet connection, the switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.  
Are you sure you want to continue? (y/n)
```

Enter `y` to continue and run the diagnostics. Enter `n` to cancel the operation.

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 11.0 to run diagnostics on management modules from the command line interface (CLI).

The `stack-port` parameter for the Summit family of switches was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

run elrp

```
run elrp <vlan_name> {ports <ports>} {interval <sec>} {retry <count>}
```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.

Syntax Description

vlan_name	Specifies a VLAN name.
ports	Specifies the set of VLAN ports for packet transmission.
sec	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 64 seconds. The default is 1 second.
count	Specifies the number of times ELRP packets must be transmitted. The range is 3 to 255 times. The default is 10 times.

Default

- `sec`—The interval between consecutive packet transmissions is 1 second.
- `count`—The number of time ELRP packets must be transmitted is 10.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client prints a log message to the console. There is no need to send a trap to the SNMP manager for non-periodic requests.



NOTE

This command is compatible with Extreme Networks switches running only the ExtremeXOS software. If your network contains switches running ExtremeXOS and switches running ExtremeWare, use the `configure elrp-client one-shot <vlan_name> ports [<ports> | all] interval <sec> retry <count> [log | print | print-and-log]` command to perform one-time ELRP packet transmission.

If you do not specify the optional `interval` or `retry` parameters, the default values are used.

Use the `configure elrp-client periodic` command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following command starts one-time, non-periodic ELRP packet transmission on the VLAN *green* using the default interval and packet transmission:

```
run elrp green
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

save debug tracefiles memorycard

```
save debug tracefiles memorycard
```

Description

Copies debug information to the external memory card.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines



NOTE

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

Use this command to copy debug information to the specified memory card. The debug information includes log files and trace files.

Progress messages are displayed that indicate the file being copied and when the copying is finished.

Beginning with ExtremeXOS 11.6, you can use the `upload debug [<hostname> | <ipaddress>] {{vr} <vrname>}` command to copy debug information to a network TFTP server.

BlackDiamond 10808 Switch Only. In ExtremeXOS 11.0, use the `upload debug-info memorycard` command to copy debug information to the external memory card.

Example

The following command copies debug information to the installed external memory card:

```
save debug tracefiles memorycard
```

History

This command was first available in ExtremeXOS 11.0.

The syntax for this command was modified in ExtremeXOS 11.1 from `upload debug-info memorycard` to `save debug tracefiles memorycard`.

Platform Availability

This command is available only on modular switches.

show debug

show debug

Description

This command displays the status of writing core dump files to the specified memory card.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

By default, the switch does not write core dump files to the memory card. Use this command to verify if you are writing core dump files to the memory card. By default, the switch does not write core dump files to the memory card.

All Switches. To configure the switch to write core dump files to the internal memory card, use the `configure debug core-dumps internal-memory` command.

Modular Switches Only. To configure a modular switch to write core dump files to the external memory card, use the `configure debug core-dumps memorycard` command.

Example

The following example shows if the switch is sending core dump files to the specified memory card:

```
show debug
```

Output for All Switches. The following sample output shows that the switch is sending core dump files to the internal memory card:

```
Debug Settings:
  Core dumps: Enabled (internal-memory)
```

The following sample output shows that the switch is not sending core dump files to the specified memory card (this is the default behavior):

```
Debug Settings:
  Core dumps: Disabled
```

Modular Switches Only. The following sample output shows that the switch is sending core dump files to the external memory card:

```
Debug Settings:
  Core dumps: Enabled (memorycard)
```

History

This command was first available in ExtremeXOS 11.1.

Support for the internal memory card was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show diagnostics

```
show diagnostics {slot [<slot> | A | B]}
```

On a stack:

```
show diagnostics {slot [<slotnumber>]}
```

Description

Displays the status of the last diagnostic test run on the switch.

Syntax Description

slot	Specifies which I/O module to display diagnostic status information on. NOTE: This parameter is available only on modular switches.
A B	Specifies which MSM to display diagnostic status information on: • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.
slot number	Specifies the slot number of the node on which you need to run the diagnostics. Values can be from 1 to 8. Note: This parameter is available only on a stackable switch in stack mode.

Default

N/A.

Usage Guidelines

Use this command to display information from the last diagnostic test run on the switch.

To run diagnostics on a switch in a stack, you need to remove the node from the stacking mode, restart the node, and then run the diagnostics locally on the node. You can see the latest diagnostics results from the Master node of a stack after the node has rejoined the stack.

Output on the BlackDiamond 10808 and BlackDiamond 12800 Series Switches. The switch displays the following diagnostic information:

- Type of slot—The slot where the diagnostic test was run: I/O or MSM.
- Slot number—The number of the slot where the diagnostic test was run.
- Date the test was run—The day, month, date, time, and year the test was run.
- Temperature—The temperature of the module, in celsius, when the test was run.

- Test data—More detailed information about the test, which includes the:
 - Version number of the diagnostic test run.
 - Serial number of the module where the test was run.
 - Type of tests run and if they passed or failed (CPU, register, memory, and system).
- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic test has failed.

Output on the BlackDiamond 8800 Series Switch. The switch displays the following diagnostic information:

If you have run diagnostics, a brief summary of the overall diagnostic test is displayed. Options are:

- Diagnostics Pass—The diagnostic test has passed.
- Diagnostics Fail—One or more diagnostic test has failed.
- Diagnostics Interrupted—The diagnostic test was interrupted on the I/O module due to initiating an MSM failover.

If you have never run diagnostics on a specific slot, the switch displays a message similar to the following:

```
Diagnostics never ran on slot 6
```

If you attempt to view diagnostics information for a slot that does not have a module installed, the switch displays a message similar to the following:

```
No card in slot 7
```

Output on the Summit Family of Switches. The switch displays the following diagnostic information:

If you have run diagnostics, information includes the:

- Date the test was run—The month, date, and year the test was run.
- Test data—If the diagnostic test failed, the switch displays the name of the failed test. The switch displays a message similar to the following:

```
MAC memory test failed
```

- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic tests have failed.

If you have never run diagnostics on the switch or stack ports, the switch displays a message similar to the following:

```
Result: FAIL
      Test date run is invalid. Please run Diagnostics.
Error reading diagnostics information.
```

This message is normal and expected if you have never run diagnostics on the switch. After you run diagnostics, you should see information about the executed test.

Additional Guidelines Applicable to Modular Switches Only. If you use the `show diagnostics {slot [<slot> | A | B]}` command on a slot where diagnostics have not been run, the switch displays messages similar to the following:

```
Diagnostics never ran on slot 3
```

Or

```
Diagnostics never run
```

If you try to display diagnostic test information on a slot where no module is installed, the switch displays messages similar to the following:

```
No card in slot 3
```

Or

```
No card in slot
```

Running Diagnostics. To run diagnostics on an I/O module or an MSM installed in the BlackDiamond 10808 switch or the BlackDiamond 8800 series switch, use the following command:

```
run diagnostics [extended | normal] {slot [<slot> | A | B]}
```

To run diagnostics on all of the installed I/O modules and the primary MSM installed in the BlackDiamond 12800 series switches, use the following command:

```
run diagnostics [extended | normal]
```

To run diagnostics on the Summit family of switches, use the following command:

```
run diagnostics [extended | normal | stack-port]
```

Depending on the software version running on your switch or your switch model, additional or different diagnostic information might be displayed. For more information, see the command [run diagnostics on page 2437](#).

Example

The following command displays the results of module diagnostics for the I/O module in slot 2:

```
show diagnostics slot 2
```

The following is sample output from the BlackDiamond 10808 switch:

```
I/O SLOT
Slotnum 2
DiagResult timestamp Tue Jun 28 10:26:00 2005
Test Temperature = 29 degrees C
Diag Test Version=1.0.6.2.      S/N=0405F-00063
-----
CPU System |   Passed
-----
Register Test |   Passed
-----
Memory Test |   Passed
-----
System Test |   Passed
-----
Summary: Diagnostics Pass
```

The following is sample output from the BlackDiamond 8800 series switch:

```
Summary: Diagnostics Pass
```

The following is sample output from the BlackDiamond 12804 switch:

```
Slotnum 2
DiagResult timestamp Fri Dec  2 11:15:29 2005
Test Temperature = 41 degrees C
Diag Test Version=1.0.1.2.      S/N=05364-00058
-----
CPU System |   Passed
-----
Register Test |   Passed
-----
Memory Test |   Passed
-----
System Test |   Passed
-----
Summary: Diagnostics Pass
```

The following command displays the results of the switch diagnostics for the Summit X450 family switch:

```
show diagnostics
```

The following is sample output from this command:

```
Last Test Date: May-04-2006
Summary: Diagnostics Pass
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show elrp

show elrp

Description

Displays ELRP information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If you enter the `show elrp` command, this command displays the:

- State of ELRP (enabled/disabled).
- Total number of ELRP sessions.
- ELRP packets transmitted.
- ELRP packets received.

In addition to the summary information at the top of the display, the `show elrp` command also displays the following information:

Client	Displays the name of the ELRP client.
VLAN	Displays the name of the VLAN with ELRP enabled.
Ports	Displays the set of VLAN ports used for packet transmission.
Interval	Displays the configured interval. An interval of 3 indicates that ELRP PDUs are transmitted every 3 seconds.
Count	Lists the configured number of ELRP PDUs that are transmitted. The PDUs are transmitted at the configured interval. This method of ELRP PDU transmission is used by ESRP in the pre-master state. A count of 0 indicates continuous PDU transmission. If the Cyclic value is <code>Yes</code> , the count is always 0.
Cyclic	Indicates whether ELRP PDUs are being continuously sent. The column shows <code>Yes</code> for the master VLAN because that VLAN is continuously sending ELRP PDUs for loop detection. When a VLAN is in the pre-master state, it only sends three ELRP PDUs before changing to master or slave. During this time the column shows <code>No</code> for that VLAN.
Pkts-Xmit	Displays the number of ELRP PDUs transmitted.
Pkts-Rcvd	Displays the number of ELRP PDUs received.

Action	Displays the configured action the switch takes when ELRP messages are received back indicating a detection of a network loop or no packets are received within the specified duration. The following list describes the actions: • Print (P)—Specifies that the switch prints a message to the console. • Log (L)—Specifies that the switch sends a message to the system log file. • Trap (T)—Specifies that the switch sends a message to the SNMP manager. • Callback (C)—Specifies a callback action. If you use ELRP with another protocol (for example ESRP), ELRP uses a callback action to notify the protocol of a loop detection.

Example

The following command displays summary ELRP status information on the switch:

```
show elrp
```

The following sample output is displayed:

```
ELRP Standalone Client:      Enabled
```

```
Number of ELRP sessions:      1
Number of ELRP pkts transmitted: 3288
Number of ELRP pkts Received:  0
```

Client	Vlan	Ports	Int.	Count	Cyclic	Pkts-Xmit	Pkts-Rcvd	Action
CLI	Default	All	1	0	Yes	3288	0	LT

```
Action : (P) Print , (L) Log , (T) Trap , (C) Callback
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show forwarding configuration

```
show forwarding configuration
```

Description

Displays the hardware table settings, including the configured and current hash algorithm.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- `Configured hash`—The hash algorithm configured on the switch. After a switch reboot, the switch uses this hash algorithm.
- `Current hash`—The hash algorithm currently used by the switch.

It is possible for the values of the configured hash and the current hash to be different. For example, if you modified the hash algorithm and have not rebooted the switch, the values might be different. In this situation, the switch also displays the following message:

```
A save and reboot are required before the configured hash will take effect
```

Example

The following command displays the hardware forwarding algorithm configured on the switch:

```
show forwarding configuration
```

The following is sample output from this command:

```
Forwarding table hash algorithm
Configured hash:  crc16
Current hash:     crc16
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available only on the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches.

show tech

```
show tech {all | <area>} {detail} {logto [file]}
```

Description

Displays the output of various show commands to assist in monitoring and troubleshooting the switch; use only in conjunction with Extreme Networks Technical support.

Syntax Description

all	Indicates all available show command output to be displayed.
area	Specifies one tech support area. For example, if you want to view STP information, enter <code>stp</code> .
detail	Specifies more detailed information.
logto [file]	Instructs the switch to log the show tech output into a file located in the switch's internal memory. The default file name is <code>show_tech.log.tgz</code> .

Default

N/A.

Usage Guidelines



NOTE

Use this command only under the guidance of Extreme Networks Technical Support personnel to view your switch configurations and to troubleshoot the switch.

The `show tech` command displays the output of the following `show` commands, among others:

- `show bootprelay`
- `show configuration`
- `show dhcp-client state`
- `show diagnostics`
- `show memory`
- `show odometers`
- `show policy`
- `show port rxerror`
- `show port txerror`
- `show power`
- `show power budget`
- `show power controller`
- `show process`

- `show radius`
- `show session`
- `show switch`
- `show tacacs`
- `show version`
- `show vlan`

Information about the following areas is also displayed, among others:

- `aaa`
- `bootp`
- `cli`
- `stp`

If you enter the `detail` keyword, the following show output is displayed, among others:

- `show log`
- `show log configuration`
- `show log counters all`
- `show process detail`

This information can be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch, the configurations running on your switch, and the type of switch you have, additional or different `show` command and configuration output may be displayed.

Example

The following command displays the show command output on the switch:

```
show tech
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

top

top

Description

Displays real-time CPU utilization information by process.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the percentage of CPU processing devoted to each process, sampled every 5 seconds.

You can change the display by typing a character while the display is active. [Table 33](#) displays the supported commands.

Table 33: TOP interactive command display options

Key	Action
P	Sort process list by CPU utilization
T	Sort process list by time usage
N	Sort process list by number (process ID)
M	Sort process list by memory usage
q [Ctrl] + c	Exit the top program

For more detailed information about the top command including display options, command fields, and command usage, please refer to your UNIX documentation.

Example

The following command displays the real-time CPU utilization information by process:

```
top
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure elrp-client

```
unconfigure elrp-client <vlan_name>
```

Description

Disables a pending one-shot or periodic ELRP request for the specified VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.

Default

None.

Usage Guidelines

This command disables a pending one-shot or periodic ELRP request for the specified VLAN.

To start one-time, non-periodic ELRP packet transmission on specified ports of a VLAN using a particular count and interval, use one of the following commands:

- `configure elrp-client one-shot <vlan_name> ports [<ports> | all] interval <sec> retry <count> [log | print | print-and-log]`—(This command is backward compatible with Extreme Networks switches running the ExtremeWare software.)
- `run elrp <vlan_name> {ports <ports>} {interval <sec>} {retry <count>}`

To configure periodic transmission of ELRP packets, use the `configure elrp-client periodic` command.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following command disables a pending ELRP request on VLAN *elrp1*:

```
unconfigure elrp-client elrp1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

upload debug

```
upload debug [<hostname> | <ipaddress>] [{vr} <vrname>]
```

Description

Uploads debug information files to a tftp server. On a platform that has both primary and backup MSMs, debug information files are uploaded from both the backup and primary MSM's.

Syntax Description

hostname	Specifies the host name of the TFTP server to which the debug files will be uploaded to.
ipaddress	Specifies the IP address of the TFTP server to which the debug files will be uploaded to.
vrname	Specifies the name of the virtual router.

Default

By default, the virtual router VR-Mgmt will be used.

Usage Guidelines



NOTE

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

Use this command to copy, upload debug information (for example, core, trace, show tech, configuration, and policy files) to the specified TFTP server.

Progress messages are displayed that indicate the file being copied and when the copying is finished. Depending on your platform, the switch displays a message similar to the following:

```
The following files on have been uploaded:
Tarball Name: TechPubsLab_C_09271428.tgz
./primary.cfg
```

You can also use this command in conjunction with the [show tech](#) command. Prior to uploading debug information files, the switch prompts you with the following message to run the [show tech](#) command with the logto file option:

```
Do you want to run show tech logto file first? (y/n)
```

Enter [y](#) to run the [show tech](#) command before uploading debug information. If you enter [y](#), the `show_tech.log.tgz` file is included during the upload. Enter [n](#) to upload debug information without running the [show tech](#) command.

After you upload the debug information, you should see a compressed TAR file, which contains the debug information.

Example

The following command uploads debug files to a network TFTP server:

```
upload debug 10.10.10.10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

The Converged Network Analyzer (CNA) Agent is part of the CNA software used to test network conditions from Avaya Inc.® You use the CNA Agent *only* if your network includes an Avaya solution that uses CNA.

The entire CNA software package consists of multiple parts. The user interface is a combination of a Java applet hosted from the CNA Server and a Command Line Interface (CLI). You obtain all parts (except the CNA Agent) from your Avaya representative, along with the accompanying documentation. You configure and manage the CNA Agent using the CLI.

The CNA Agent is a downloadable software module that you use when running the Avaya CNA software.

**NOTE**

You must download and install the Secure Shell (SSH) software module prior to downloading and installing the CNA Agent software module.

Using the CNA software, the CNA Agent runs the requested tests and returns the test results. The CNA Agent runs the following tests as directed by the CNA Server:

- Traceroute
- RTP
- Ping
- TCPconnect
- Merge

You enable the software and configure the CNA Agent to communicate with the CNA Server, to clear the test counters, and to display connection status and test results.

clear cna-testplug counters

```
clear cna-testplug counters
```

Description

Clears all counters maintained by the CNA Agent and resets the counters to zero.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You can also use the `clear counters` command to reset the internal counters for the CNA Agent and return them to 0.

Example

The following command clears all the counters on the CNA Agent and returns the values to zero:

```
clear cna-testplug counters
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure cna-testplug scheduler ipaddress

```
configure cna-testplug scheduler ipaddress <ip_address>
```

Description

Configures the IP address of the CNA Server using open Secure Socket Layer (SSL); you enter the IP address of the CNA server.

Syntax

ip_address	Specifies the IP address of the CNA Server that communicates with the CNA Agent to schedule tests and receive the results.

Default

N/A.

Usage Guidelines

Use this command to configure the CNA Agent with the IP address of the CNA Server. The CNA Server requests the timing and type of networking testing, and the CNA Agent runs the tests.



NOTE

You use the CNA Agent only if you are running the Avaya CNA solution; you must have other pieces of the CNA (available from Avaya) to run these tests.

You enter the IP address of the CNA server.

This command sets up the encryption key that is subsequently used for all communication between the CNA Agent and the CNA Server.



NOTE

You must have previously installed the Secure Shell (SSH) downloadable software module, which contains SSL, to use the CNA Agent software.

Example

The following command enters the CNA Server's IP address to 10.6.13.116; the CNA Agent uses this IP address to communicate with the CNA Server:

```
configure cna-testplug scheduler ipaddress 10.6.13.116
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure cna-testplug vlan

```
configure cna-testplug vlan <vlan_name>
```

Description

Configures the CNA Agent (test plug) to an interface. By default, the CNA Agent is bound to the Default VLAN.

Syntax

valn_name	Specifies which interface IP address the CNA Server uses to communicate with the CNA Agent.

Default

Default VLAN.

Usage Guidelines

The interface IP address is specified when setting up the SSL connection with the CNA Server (when you issue the `configure cna-testplug scheduler ipaddress` command. The CNA Server attempts to establish the socket connection on the interface specified in the `configure cna-testplug vlan` interface to conduct the actual tests.

The system uses the primary IP address if the VLAN has more than one IP address. By default, the CNA Server uses the Default VLAN.



NOTE

Extreme Networks recommends that you put IP telephones on the same virtual router.

Example

The following command instructs the CNA Server to use the interface associated with VLAN gateway to conduct tests:

```
configure cna-testplug vlan gateway
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable cna-testplug

```
disable cna-testplug
```

Description

Disables the CNA Agent.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

You must enable the CNA Agent before it can run the network tests requested by the CNA Server.

Example

The following command disables the CNA Agent:

```
disable cna-testplug
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable cna-testplug

```
enable cna-testplug
```

Description

Enables the CNA Agent (test plug). Once enabled, the CNA Agent coordinates with the CNA Server to test the network for throughput.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If you previously configured the IP address for the CNA Server and the VLAN interface, the CNA Agent immediately registers with the CNA Server upon being enabled and begins running the requested tests.

Example

The following command enables the CNA Agent:

```
enable cna-testplug
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show cna-testplug

```
show cna-testplug
```

Description

Displays the statistics and connection status with the CNA software from Avaya. The display includes configured CNA Agent (test plug) and CNA Server (scheduler) connections and the number of tests conducted on each connection.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information:

- Hardware Name—The name of the Extreme Networks device running the CNA Agent (test plug)
- Firmware version—The version of ExtremeXOS firmware running on the device
- Interface VLAN—VLAN (and virtual router) interface the CNA Server uses to schedule and run the tests
- IP address on the interface—IP address of the interface the CNA Server uses; the connection on which the CNA Agent tests and sends results
- CNA Test plug version—Version of CNA Agent (test plug) software, in the following format: MajorRev.MinorRev.Build
- Interface version with Scheduler—Shows compatibility between the CNA Agent (test plug) and the CNA Server (scheduler) software
- Enabled or Disabled
- Status—Connection status to the CNA Server:
 - Registered
 - Unregistered
- Errors—Number of errors in tests of connectivity
- Total tests received—Total test requests received by the CNA Agent
- Scheduler (SBC)—IP address and port number on which the CNA Server communicates with the CNA Agent; the connection on which the CNA Agent listens

- Listening ports
 - Test requests(from ANS)—Number of tests that the CNA Agent was requested to run by the Adaptive Networking Software (ANS) on the CNA Server
 - RTP test requests(from test plugs)—Number of RTP streams that the CNA Agent initiates
 - RTP and Traceroute responses(from test plugs)—Number of responses by the CNA Agent to requests for RTP and Traceroute tests
- Last Test—Last test that CNA Agent performed
- Result Last Test—Results of the last test, which the CNA Agent sends to the CNA Server
- Test—Tests run by the CNA Agent
- Count—Number of tests successfully run

Example

The following command displays CNA Agent statistics and connection status on the BlackDiamond 10808 switch:

```
show cna-testplug
```

```
HW Name: BlackDiamond 10808
Firmware Version: 11.3.0.4
Interface VLAN: "Default" on Virtual router "VR-Default"
IP address of the Interface:      10.203.135.61
CNA Testplug version: 3.0.1
Interface version with Scheduler:17
Admin: Enabled
Status: Registered
Errors: 29846
Total tests received: 58338
Scheduler (SBC): 10.203.128.123: 20000:
Listening ports
--Test requests(from ANS): 50000
--RTP test requests(from test plugs): 50001
--RTP and Traceroute responses(from test plugs): 50015
Last Test: Ping to 10.203.135.61
Result Last Test: ping results: rtt: 2.292000, loss: 0.000000
Results:


| TEST       | COUNT  | FAILED |
|------------|--------|--------|
| Traceroute | 65489  | 0      |
| RTP        | 51403  | 0      |
| Ping       | 101549 | 0      |
| Tcpconnect | 36455  | 0      |
| Merge      | 50     | 0      |


```

The following command displays CNA Agent statistics and connection status on the Summit X450 series switch:

```
show cna
```

```
HW Name: SummitX450-24t
Firmware version: 11.3.0.11
Interface VLAN: "Default" on Virtual router "VR-Default"
IP address of the Interface: 10.203.128.126
```

```

CNA Test plug version:      3.0.2
CNA Interface version:      17
Admin:                      Enabled
Status:                     Registered
Errors:                     22
Total tests received:       23529
Scheduler (SBC):            10.203.128.127
Listening ports
  --Test requests(from ANS):      50000
  --RTP test requests(from test plugs):  50001
  --RTP and Traceroute responses(from test plugs): 50016
Last Test: RTP to 10.203.128.124
Result Last Test: delay: 0.136375, jitter: 0.007525, loss: 0.000000(in milli secs)

```

Test Results:

Name	COUNT	FAILED
----	-----	-----
Traceroute	74898	0
RTP	145146	22
Ping	0	0
Tcpconnect	0	0
Merge	93	0



NOTE

Adaptive Networking Software (ANS) runs on the CNA Server.

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

Index of Commands

C

- check policy, 924
- check policy attribute, 925
- clear access-list counter, 934
- clear access-list meter, 936
- clear account lockout, 56
- clear bgp flap-statistics, 2086
- clear bgp neighbor counters, 2085
- clear bgp neighbor flap-statistics, 2086
- clear cna-testplug counters, 2462
- clear counters, 649
- clear counters edp, 323
- clear counters mpls ldp, 1581
- clear counters ports, 321
- clear counters stp, 1422
- clear counters traffic queue, 977
- clear counters vrrp, 1551
- clear cpu-monitoring, 270
- clear eaps counters, 1357
- clear elrp counters, 2421
- clear elsm counters, 651
- clear elsm ports auto-restart, 650
- clear elsm ports counters, 651
- clear esrp counters, 1490
- clear esrp neighbor, 1492
- clear esrp sticky, 1493
- clear fdb, 898
- clear fdb vlpls, 1584
- clear igmp group, 2202
- clear igmp snooping, 2203
- clear inline-power stats ports, 596
- clear iparp, 1743
- clear ip-security anomaly-protection notify cache, 1070
- clear ip-security dhcp-snooping entries, 1071
- clear ip-security source-ip-lockdown entries ports, 1072
- clear lacp counters, 322
- clear license-info, 57
- clear log, 652
- clear log counters, 653
- clear mac-binding, 820
- clear mld group, 2284
- clear mld snooping, 2285
- clear msdp counters, 2310
- clear msdp sa-cache, 2312
- clear netlogin state, 1246
- clear netlogin state mac-address, 1247
- clear ospf counters, 1961
- clear ospfv3 counters, 2028
- clear pim cache, 2204
- clear rip counters, 1890
- clear ripng counters, 1928
- clear session, 58
- clear slot, 324
- clear sys-recovery-level, 655
- clear vlan dhcp-address-allocation, 1073
- config inline-power priority ports, 606
- configure access-list, 937
- configure access-list add, 939
- configure access-list delete, 941, 942
- configure account, 59
- configure account encrypted, 61
- configure account password-policy char-validation, 63
- configure account password-policy history, 65
- configure account password-policy lockout-on-login-failures, 67
- configure account password-policy max-age, 69
- configure account password-policy min-length, 71
- configure banner, 73
- configure bgp add aggregate-address, 2088
- configure bgp add confederation-peer sub-AS-number, 2090
- configure bgp add network, 2091
- configure bgp AS-number, 2092
- configure bgp cluster-id, 2093
- configure bgp confederation-id, 2094
- configure bgp delete aggregate-address, 2095
- configure bgp delete confederation-peer sub-AS-number, 2096
- configure bgp delete network, 2097
- configure bgp export shutdown-priority, 2098
- configure bgp import-policy, 2100
- configure bgp local-preference, 2101
- configure bgp med, 2102
- configure bgp neighbor dampening, 2103
- configure bgp neighbor description, 2105
- configure bgp neighbor maximum-prefix, 2106
- configure bgp neighbor next-hop-self, 2108
- configure bgp neighbor no-dampening, 2109
- configure bgp neighbor password, 2110
- configure bgp neighbor peer-group, 2112
- configure bgp neighbor route-policy, 2114

configure bgp neighbor route-reflector-client, 2115
 configure bgp neighbor send-community, 2116
 configure bgp neighbor shutdown-priority, 2117
 configure bgp neighbor soft-reset, 2118
 configure bgp neighbor source-interface, 2120
 configure bgp neighbor timer, 2121
 configure bgp neighbor weight, 2122
 configure bgp peer-group dampening, 2123
 configure bgp peer-group maximum-prefix, 2125
 configure bgp peer-group next-hop-self, 2127
 configure bgp peer-group no-dampening, 2128
 configure bgp peer-group password, 2131
 configure bgp peer-group remote-AS-number, 2132
 configure bgp peer-group route-policy, 2133
 configure bgp peer-group route-reflector-client, 2129
 configure bgp peer-group send-community, 2130
 configure bgp peer-group soft-reset, 2134
 configure bgp peer-group source-interface, 2136
 configure bgp peer-group timer, 2138
 configure bgp peer-group weight, 2140
 configure bgp restart, 2141
 configure bgp restart address-family, 2142
 configure bgp restart restart-time, 2143
 configure bgp restart stale-route-time, 2144
 configure bgp restart update-delay, 2145
 configure bgp routerid, 2146
 configure bgp soft-reconfiguration, 2147
 configure bootprelay add, 1744
 configure bootprelay delete, 1745
 configure bootprelay dhcp-agent information check, 1746
 configure bootprelay dhcp-agent information option, 1747
 configure bootprelay dhcp-agent information policy, 1748
 configure bvlan add ports tagged, 821
 configure bvlan add svlan, 822
 configure bvlan delete ports, 823
 configure bvlan delete svlan, 824
 configure bvlan ethertype, 825
 configure bvlan tag, 826
 configure cfm domain add association integer, 510
 configure cfm domain add association string, 511
 configure cfm domain add association vpn-id oui index, 512
 configure cfm domain association add end-point mepid, 513
 configure cfm domain association add filter-function, 515
 configure cfm domain association add intermediate-point, 517
 configure cfm domain association delete end-point, 519
 configure cfm domain association delete filter-function, 520
 configure cfm domain association delete intermediate-point, 521
 configure cfm domain association end-point mepid, 522
 configure cfm domain association end-point transmit-interval, 523
 configure cfm domain delete association, 525
 configure cli max-failed-logins, 75
 configure cli max-sessions, 74
 configure cna-testplug scheduler ipaddress, 2463
 configure cna-testplug vlan, 2465
 configure debug coredumps, 2422
 configure diffserv examination code-point, 978
 configure diffserv replacement, 980
 configure dns-client add, 76
 configure dns-client default-domain, 78
 configure dns-client delete, 79
 configure dos-protect acl-expire, 1074
 configure dos-protect interval, 1075
 configure dos-protect trusted-ports, 1076
 configure dos-protect type l3-protect alert-threshold, 1077
 configure dos-protect type l3-protect notify-threshold, 1078
 configure dot1p type, 982
 configure eaps add control vlan, 1358
 configure eaps add protected vlan, 1360
 configure eaps config-warnings off, 1362
 configure eaps config-warnings on, 1364
 configure eaps delete control vlan, 1365
 configure eaps delete protect vlan, 1366
 configure eaps failtime, 1368
 configure eaps failtime expiry-action, 1370
 configure eaps fast-convergence, 1371
 configure eaps hellotime, 1373
 configure eaps mode, 1374
 configure eaps name, 1376
 configure eaps port, 1377
 configure eaps shared-port link-id, 1379
 configure eaps shared-port mode, 1380
 configure eaps shared-port segment-timeout expiry-action, 1381
 configure edp advertisement-interval, 325
 configure elrp-client one-shot, 2424
 configure elrp-client periodic, 2426
 configure elsm ports hellotime, 657
 configure elsm ports hold-threshold, 659, 661
 configure elsm ports uptimer-threshold, 661

configure esrp add elrp-poll ports, 1494
 configure esrp add master, 1495
 configure esrp add member, 1496
 configure esrp add track-environment, 1497
 configure esrp add track-iproute, 1498
 configure esrp add track-ping, 1499
 configure esrp add track-vlan, 1500, 1501, 1502
 configure esrp delete elrp-poll ports, 1503
 configure esrp delete master, 1504
 configure esrp delete member, 1505
 configure esrp delete track-environment, 1506
 configure esrp delete track-iproute, 1507
 configure esrp delete track-ping, 1508
 configure esrp delete track-vlan, 1509
 configure esrp domain-id, 1510
 configure esrp election-policy, 1511
 configure esrp elrp-master-poll disable, 1515
 configure esrp elrp-master-poll enable, 1516
 configure esrp elrp-premaster-poll disable, 1518
 configure esrp elrp-premaster-poll enable, 1519
 configure esrp group, 1521
 configure esrp mode, 1523
 configure esrp name, 1525
 configure esrp ports mode, 1526
 configure esrp ports no-restart, 1527
 configure esrp ports restart, 1528
 configure esrp ports weight, 1529
 configure esrp priority, 1530
 configure esrp timer hello, 1531
 configure esrp timer neighbor, 1532
 configure esrp timer neutral, 1534
 configure esrp timer premaster, 1535
 configure esrp timer restart, 1536
 configure failsafe-account, 81
 configure fdb agingtime, 899
 configure firmware, 2362
 configure flood-group extended add, 874
 configure flood-group extended delete, 875
 configure flood-group virtual add, 876
 configure flood-group virtual delete, 878
 configure forwarding hash-algorithm, 2428
 configure idletimeout, 83
 configure igmp, 2205
 configure igmp snooping flood-list, 2207
 configure igmp snooping leave-timeout, 2209
 configure igmp snooping timer, 2210
 configure igmp snooping vlan ports add static group, 2211
 configure igmp snooping vlan ports add static router, 2214
 configure igmp snooping vlan ports delete static group, 2213
 configure igmp snooping vlan ports delete static router, 2215
 configure igmp snooping vlan ports filter, 2216
 configure inline-power budget slot, 598
 configure inline-power disconnect-precedence, 600
 configure inline-power label ports, 602
 configure inline-power operator-limit ports, 604
 configure inline-power usage-threshold, 607
 configure iparp add, 1749
 configure iparp add proxy, 1750
 configure iparp delete, 1752
 configure iparp delete proxy, 1753
 configure iparp max_entries, 1754
 configure iparp max_pending_entries, 1755
 configure iparp max_proxy_entries, 1756
 configure iparp timeout, 1757
 configure ip-mtu vlan, 326
 configure iproute add, 1759, 1843
 configure iproute add blackhole, 1760, 1844
 configure iproute add blackhole default, 1761
 configure iproute add blackhole ipv6 default, 1845
 configure iproute add default, 1762, 1846
 configure iproute delete, 1763, 1847
 configure iproute delete blackhole, 1764, 1848
 configure iproute delete blackhole default, 1765
 configure iproute delete blackhole ipv6 default, 1849
 configure iproute delete default, 1766, 1850
 configure iproute ipv6 priority, 1851
 configure iproute priority, 1767
 configure iproute sharing max-gateways, 1769
 configure ip-security anomaly-protection icmp ipv4-max-size, 1079
 configure ip-security anomaly-protection icmp ipv6-max-size, 1080
 Configure ip-security anomaly-protection notify cache, 1081
 Configure ip-security anomaly-protection notify rate limit, 1082
 configure ip-security anomaly-protection notify rate window, 1083
 configure ip-security anomaly-protection notify trigger off, 1084
 configure ip-security anomaly-protection notify trigger on, 1085
 configure ip-security anomaly-protection tcp, 1086
 configure irdp, 1770
 configure jumbo-frame size, 328
 configure lacp member-port priority, 330
 configure lldp med fast-start repeat-count, 544
 configure lldp ports management-address, 545
 configure lldp ports port-description, 546
 configure lldp ports system-capabilities, 547

configure lldp ports system-description, 548
 configure lldp ports system-name, 549
 configure lldp ports vendor-specific avaya-extreme call-server, 550
 configure lldp ports vendor-specific avaya-extreme dot1q-framing, 551
 configure lldp ports vendor-specific avaya-extreme file-server, 553
 configure lldp ports vendor-specific avaya-extreme poe-conservation-request, 554
 configure lldp ports vendor-specific dot1 port-protocol-vlan-ID, 557
 configure lldp ports vendor-specific dot1 port-vlan-id, 556
 configure lldp ports vendor-specific dot1 vlan-name, 559
 configure lldp ports vendor-specific dot3 link-aggregation, 561
 configure lldp ports vendor-specific dot3 mac-phy, 562
 configure lldp ports vendor-specific dot3 max-frame-size, 563
 configure lldp ports vendor-specific dot3 power-via-mdi, 564
 configure lldp ports vendor-specific med capabilities, 566
 configure lldp ports vendor-specific med location-identification, 568
 configure lldp ports vendor-specific med policy application, 570
 configure lldp ports vendor-specific med power-via-mdi, 573
 configure lldp reinitialize-delay, 575
 configure lldp snmp-notification-interval, 576
 configure lldp transmit-delay, 577
 configure lldp transmit-hold, 578
 configure lldp transmit-interval, 579
 configure log display, 662
 configure log filter events, 664
 configure log filter events match, 667
 configure log filter events strict-match, 667
 configure log target filter, 671
 configure log target format, 674
 configure log target match, 678
 configure log target severity, 680
 configure log target syslog, 682
 configure mac-lockdown-timeout ports aging-time, 1087, 1147
 configure meter, 984
 configure meter peak-rate, 986, 988
 configure mirror add anomaly, 331
 configure mirroring add, 332
 configure mirroring delete, 337
 configure mirroring mode, 338
 configure mld, 2286
 configure mld snooping flood-list, 2291
 configure mld snooping leave-timeout, 2293
 configure mld snooping timer, 2294
 configure mld snooping vlan ports add static group, 2287
 configure mld snooping vlan ports add static router, 2289
 configure mld snooping vlan ports delete static group, 2288
 configure mld snooping vlan ports delete static router, 2290
 configure mpls, 1601
 configure mpls add vlan, 1592
 configure mpls delete vlan, 1593
 configure mpls ldp advertise, 1596
 configure mpls ldp loop-detection, 1598
 configure mpls ldp timers, 1599
 configure msdp default-peer, 2313
 configure msdp max-rejected-cache, 2315
 configure msdp mesh-group, 2316
 configure msdp no-default-peer, 2318
 configure msdp originator-id, 2319
 configure msdp password, 2320
 configure msdp peer description, 2322
 configure msdp sa-cache-server, 2323
 configure msdp sa-filter, 2325
 configure msdp sa-limit, 2327
 configure msdp source-interface, 2328
 configure msdp timer keep-alive, 2330
 configure msdp ttl-threshold, 2331
 configure mstp format, 1423
 configure mstp region, 1424
 configure mstp revision, 1426
 configure mvr add receiver, 2221
 configure mvr add vlan, 2222
 configure mvr delete receiver, 2223
 configure mvr delete vlan, 2224
 configure mvr mvr-address, 2225
 configure mvr static group, 2227
 configure neighbor-discovery cache, 1842
 configure neighbor-discovery cache add, 1853
 configure neighbor-discovery cache delete, 1854
 configure netlogin add mac-list, 1248
 configure netlogin agingtime, 1250
 configure netlogin authentication database order, 1251
 configure netlogin banner, 1252
 configure netlogin base-url, 1253
 configure netlogin delete mac-list, 1254
 configure netlogin dot1x eapol-transmit-version, 1255
 configure netlogin dot1x guest-vlan, 1256
 configure netlogin dot1x timers, 1258

configure netlogin dynamic-vlan, 1260
 configure netlogin dynamic-vlan uplink-ports, 1261, 1262
 configure netlogin move-fail-action, 1266
 configure netlogin port no-restart, 1270
 configure netlogin port restart, 1271
 configure netlogin redirect-page, 1272
 configure netlogin vlan, 1273
 configure node priority, 128
 configure ospf add virtual-link, 1966
 configure ospf add vlan area, 1967
 configure ospf add vlan area link-type, 1968
 configure ospf area add range, 1971
 configure ospf area delete range, 1972
 configure ospf area external-filter, 1969
 configure ospf area interarea-filter, 1970
 configure ospf area normal, 1973
 configure ospf area nssa stub-default-cost, 1974
 configure ospf area stub stub-default-cost, 1975
 configure ospf area timer, 1976
 configure ospf ase-limit, 1978
 configure ospf ase-summary add cost, 1979
 configure ospf ase-summary delete, 1980
 configure ospf authentication, 1964
 configure ospf cost, 1962
 configure ospf delete virtual-link, 1981
 configure ospf delete vlan, 1982
 configure ospf import-policy, 1983
 configure ospf lsa-batch-interval, 1984
 configure ospf metric-table, 1985
 configure ospf priority, 1963
 configure ospf restart, 1986
 configure ospf restart grace-period, 1987
 configure ospf restart-helper, 1988
 configure ospf routerid, 1990
 configure ospf spf-hold-time, 1991
 configure ospf timer, 1976, 1992, 1997, 2058
 configure ospf virtual-link timer, 1992
 configure ospf vlan area, 1994
 configure ospf vlan neighbor add, 1995
 configure ospf vlan neighbor delete, 1996
 configure ospf vlan timer, 1997
 configure ospfv3 add interface, 2030
 configure ospfv3 add tunnel, 2030
 configure ospfv3 add tunnel all, 2032
 configure ospfv3 add virtual-link, 2034
 configure ospfv3 add vlan, 2030
 configure ospfv3 add vlan all, 2032
 configure ospfv3 area add range, 2035
 configure ospfv3 area cost, 2036
 configure ospfv3 area delete range, 2037
 configure ospfv3 area external-filter, 2038
 configure ospfv3 area interarea-filter, 2040
 configure ospfv3 area normal, 2042
 configure ospfv3 area priority, 2043
 configure ospfv3 area stub, 2044
 configure ospfv3 area timer, 2045
 configure ospfv3 cost, 2036, 2051
 configure ospfv3 delete interface, 2047
 configure ospfv3 delete virtual-link, 2048
 configure ospfv3 import-policy, 2049
 configure ospfv3 interface area, 2050
 configure ospfv3 interface cost, 2051
 configure ospfv3 interface timer, 2053
 configure ospfv3 metric-table, 2055
 configure ospfv3 priority, 2043, 2052
 configure ospfv3 routerid, 2056
 configure ospfv3 spf-hold-time, 2057
 configure ospfv3 timer, 2045, 2053
 configure ospfv3 tunnel cost, 2051
 configure ospfv3 virtual-link timer, 2058
 configure ospfv3 vlan area, 2050
 configure ospfv3 vlan cost, 2051
 configure ospfv3 vlan timer, 2053
 configure ospv3 tunnel area, 2050
 configure pim add vlan, 2229
 configure pim border, 2332
 configure pim cbsr, 2231
 configure pim crp static, 2232
 configure pim crp timer, 2234
 configure pim crp vlan, 2235
 configure pim delete vlan, 2237
 configure pim register-checksum-to, 2240
 configure pim register-rate-limit-interval, 2238
 configure pim register-suppress-interval register-probe-interval, 2239
 configure pim spt-threshold, 2241
 configure pim ssm range, 2242
 configure pim timer vlan, 2244
 configure pim vlan trusted-gateway, 2245
 configure port ethertype, 880
 configure ports auto off, 339
 configure ports auto on, 341
 configure ports auto-polarity, 342
 configure ports display-string, 344
 configure ports limit-learning, 1088
 configure ports monitor vlan, 684
 configure ports preferred-medium, 345
 configure ports qosprofile, 989
 configure ports rate-limit egress, 991
 configure ports rate-limit egress aggregate, 993
 configure ports rate-limit flood, 347
 configure ports rate-limit map cos high, 995
 configure ports rate-limit packet byte-adjustment, 998
 configure ports redundant, 348
 configure ports wan-phy clocking, 350
 configure ports wan-phy framing, 351

configure ports wan-phy loopback, 352
 configure ports wan-phy trace-path, 353
 configure ports wan-phy trace-section, 354
 configure power supply, 130, 131
 configure protocol add, 827
 configure protocol delete, 829
 configure qosprofile ingress ports, 1000
 configure qosprofile ports, 1002
 configure qosprofile weight, 1005
 configure qosscheduler, 1007
 configure radius server client-ip, 1091
 configure radius shared-secret, 1093
 configure radius timeout, 1095
 configure radius-accounting server client-ip, 1096
 configure radius-accounting shared-secret, 1098
 configure radius-accounting timeout, 1100
 configure rip add vlan, 1891
 configure rip delete vlan, 1892
 configure rip garbagetime, 1893
 configure rip import-policy, 1894
 configure rip routetimeout, 1895
 configure rip rxmode, 1896
 configure rip txmode, 1897
 configure rip updatetime, 1898
 configure rip vlan cost, 1899
 configure rip vlan route-policy, 1900
 configure rip vlan rxmode, 1896
 configure rip vlan trusted-gateway, 1901
 configure rip vlan txmode, 1897
 configure ripng add, 1929
 configure ripng add tunnel, 1929
 configure ripng add vlan, 1929
 configure ripng cost, 1935
 configure ripng delete, 1930
 configure ripng delete tunnel, 1930
 configure ripng delete vlan, 1930
 configure ripng garbagetime, 1931
 configure ripng import-policy, 1932
 configure ripng route-policy, 1936
 configure ripng routetimeout, 1933
 configure ripng trusted-gateway, 1938
 configure ripng updatetime, 1934
 configure router-discovery add prefix, 1855
 configure router-discovery default-lifetime, 1857
 configure router-discovery delete prefix, 1856
 configure router-discovery hop-limit, 1858
 configure router-discovery link-mtu, 1859
 configure router-discovery managed-config-flag, 1860
 configure router-discovery max-interval, 1861
 configure router-discovery min-interval, 1862
 configure router-discovery other-config-flag, 1863
 configure router-discovery reachable-time, 1864
 configure router-discovery retransmit-time, 1865
 configure router-discovery set prefix, 1866
 configure safe-default-script, 84
 configure sflow agent ipaddress, 686
 configure sflow collector ipaddress, 687
 configure sflow max-cpu-sample-limit, 689
 configure sflow poll-interval, 690
 configure sflow ports sample-rate, 691
 configure sflow sample-rate, 693
 configure sharing add ports, 355
 configure sharing address-based, 358
 configure sharing delete ports, 360
 configure sharing lacp activity-mode, 361
 configure sharing lacp defaulted-state-action, 363
 configure sharing lacp system-priority, 365
 configure sharing lacp timeout, 367
 configure slot module, 368
 configure snmp add community, 134
 configure snmp add trapreceiver, 136
 configure snmp delete community, 138
 configure snmp delete trapreceiver, 140
 configure snmp syscontact, 141
 configure snmp syslocation, 142
 configure snmp sysname, 143
 configure snmpv3 add access, 144
 configure snmpv3 add community, 146
 configure snmpv3 add filter, 148
 configure snmpv3 add filter-profile, 150
 configure snmpv3 add group user, 151
 configure snmpv3 add mib-view, 153
 configure snmpv3 add notify, 155
 configure snmpv3 add target-addr, 156
 configure snmpv3 add target-params, 158
 configure snmpv3 add user, 160
 configure snmpv3 add user clone-from, 162
 configure snmpv3 delete access, 163
 configure snmpv3 delete community, 165
 configure snmpv3 delete filter, 166
 configure snmpv3 delete filter-profile, 167
 configure snmpv3 delete group user, 168
 configure snmpv3 delete mib-view, 170
 configure snmpv3 delete notify, 172
 configure snmpv3 delete target-addr, 173
 configure snmpv3 delete target-params, 174
 configure snmpv3 delete user, 175
 configure snmpv3 engine-boots, 177
 configure snmpv3 engine-id, 178
 configure sntp-client, 179
 configure sntp-client update-interval, 181
 configure ssh2 key, 1101
 configure sshd2 user-key add user, 1103
 configure sshd2 user-key delete user, 1104
 configure ssl certificate pregenerated, 1105

configure ssl certificate privkeylen, 1107
 configure ssl privkey pregenerated, 1109
 configure stacking alternate-ip-address, 1316
 configure stacking easy-setup, 1318
 configure stacking license-level, 1321
 configure stacking mac-address, 1323
 configure stacking master-capability, 1325
 configure stacking priority, 1327
 configure stacking redundancy, 1329
 configure stacking slot-number, 1331
 configure stacking slot-number automatic, 1331
 configure stpd add vlan, 1427
 configure stpd default-encapsulation, 1430
 configure stpd delete vlan, 1432
 configure stpd forwarddelay, 1433
 configure stpd hellotime, 1434
 configure stpd maxage, 1435
 configure stpd max-hop-count, 1436
 configure stpd mode, 1438
 configure stpd ports cost, 1440
 configure stpd ports edge-safeguard disable, 1442
 configure stpd ports edge-safeguard enable, 1444
 configure stpd ports link-type, 1446
 configure stpd ports mode, 1448
 configure stpd ports port-priority, 1450
 configure stpd ports priority, 1452
 configure stpd priority, 1454
 configure stpd tag, 1456
 configure svlan add ports, 830
 configure svlan delete ports, 831
 configure svlan isid, 832
 configure svlan tag, 833
 configure switch ports initial-mode disabled, 370
 configure switch ports initial-mode enabled, 371
 configure sys-health-check all level, 694
 configure sys-health-check interval, 697
 configure syslog add, 711
 configure syslog delete, 713
 configure sys-recovery-level, 700
 configure sys-recovery-level slot, 702
 configure sys-recovery-level switch, 709
 configure tacacs server client-ip, 1111
 configure tacacs shared-secret, 1113
 configure tacacs timeout, 1114
 configure tacacs-accounting server client-ip, 1115
 configure tacacs-accounting shared-secret, 1117
 configure tacacs-accounting timeout, 1118
 configure telnet access-profile, 182
 configure telnet port, 184
 configure telnet vr, 186
 configure time, 86
 configure timezone, 87
 configure traffic ingress queue add egress queue, 1008
 configure traffic ingress queue add egress-queue, 1008
 configure traffic ingress queue delete egress queue, 1010
 configure traffic mode, 1011
 configure traffic queue aggregate-meter, 1012
 configure traffic queue cos, 1013
 configure trusted-servers vlan add server, 1121
 configure tunnel, 1868
 configure upm event, 468
 configure upm timer after, 469
 configure upm timer at, 470
 configure upm timer profile, 472
 configure vlan add ports, 834
 configure vlan add ports stpd, 1458
 configure vlan add secondary-ipaddress, 1771
 configure vlan delete ports, 836
 configure vlan delete secondary-ipaddress, 1772
 configure vlan dhcp-address-range, 1124
 configure vlan dhcp-lease-timer, 1125
 configure vlan dhcp-options, 1126
 configure vlan ipaddress, 837
 configure vlan name, 840
 configure vlan netlogin-lease-timer, 1274
 configure vlan protocol, 841
 configure vlan qosprofile, 1015
 configure vlan tag, 843
 configure vlan udp-profile, 1773
 configure vman add ports, 883
 configure vman delete ports, 886
 configure vman ethertype, 887, 896
 configure vman rate-limit flood-traffic designated-port, 1016
 configure vman tag, 889
 configure vpls, 1624
 configure vpls add peer, 1626
 configure vpls add service, 1627
 configure vpls delete peer, 1629
 configure vpls delete service, 1630
 configure vpls peer, 1631
 configure vpls peer mpls lsp, 1632
 configure vr add ports, 914
 configure vr add protocol, 913
 configure vr delete ports, 916
 configure vr delete protocol, 915
 configure vrrp vlan vrid add ipaddress, 1552
 configure vrrp vlan vrid add track-iproute, 1553
 configure vrrp vlan vrid add track-ping, 1554
 configure vrrp vlan vrid add track-vlan, 1555
 configure vrrp vlan vrid advertisement-interval, 1556
 configure vrrp vlan vrid authentication, 1558

configure vrrp vlan vrid delete ipaddress, 1559
 configure vrrp vlan vrid delete track-iproute, 1560
 configure vrrp vlan vrid delete track-ping, 1561
 configure vrrp vlan vrid delete track-vlan, 1562
 configure vrrp vlan vrid dont-preempt, 1563
 configure vrrp vlan vrid preempt, 1564
 configure vrrp vlan vrid priority, 1565
 cp, 272
 create access-list, 944
 create account, 91
 create bgp neighbor peer-group, 2148
 create bgp neighbor remote-AS-number, 2149
 create bgp peer-group, 2150
 create bvlan, 844
 create cfm domain dns ma-level, 527
 create cfm domain iso3166 ma-level, 528
 create cfm domain m1400 ma-level, 529
 create cfm domain mac ma-level, 530
 create cfm domain ma-level, 526
 create cfm domain vpn-id oui index ma-level, 531
 create eaps, 1383
 create eaps shared-port, 1384
 create esrp, 1537
 create fdbentry vlan ports, 900
 create flood-group extended, 881
 create flood-group virtual, 882
 create log filter, 715
 create msdp mesh-group, 2333
 create msdp peer, 2334
 create netlogin local-user, 1275
 create ospf area, 1999
 create ospfv3 area, 2060
 create protocol, 845
 create qosprofile, 1019
 create sshd2 key-file, 1127
 create sshd2 user-key, 1128
 create stpd, 1461
 create svlan, 846
 create traffic queue allow-egress-shaping strict-priority, 1021
 create traffic queue egress, 1023
 create traffic queue ingress-only, 1025
 create tunnel 6to4, 1869
 create tunnel ipv6-in-ipv4, 1870
 create upm profile, 473
 create upm timer, 474
 create virtual-router, 912
 create vlan, 847
 create vman, 890
 create vpls fec-id-type pseudo-wire, 1593, 1636
 create vrrp vlan vrid, 1566

D

delete access-list, 947, 948
 delete account, 93
 delete bgp neighbor, 2151
 delete bgp peer-group, 2152
 delete bvlan, 849
 delete cfm domain, 532
 delete eaps, 1385
 delete eaps shared-port, 1386
 delete esrp, 1539
 delete fdbentry, 902
 delete log filter, 717
 delete msdp mesh-group, 2335
 delete msdp peer, 2336
 delete netlogin local-user, 1278
 delete ospf area, 2000
 delete ospfv3 area, 2061
 delete protocol, 850
 delete qosprofile, 1027
 delete sshd2 user-key, 1129
 delete stpd, 1463
 delete svlan, 851
 delete traffic queue, 1028
 delete tunnel, 1871
 delete upm profile, 475
 delete upm timer, 476
 delete virtual-router, 917
 delete vlan, 852
 delete vman, 892
 delete vpls, 1640
 delete vrrp vlan vrid, 1567
 disable access-list permit to-cpu, 949
 disable access-list refresh blackhole, 950
 disable bgp, 1641, 2153
 disable bgp aggregation, 2154
 disable bgp always-compare-med, 2155
 disable bgp community format, 2156
 disable bgp export, 2157
 disable bgp neighbor, 2159
 disable bgp neighbor capability, 2160
 disable bgp neighbor remove-private-AS-numbers, 2162
 disable bgp neighbor soft-in-reset, 2163
 disable bgp neighbor use-ip-router-alert, 2164
 disable bgp peer-group, 2165
 disable bgp peer-group capability, 2166
 disable bgp peer-group remove-private-AS-number, 2167
 disable bgp peer-group soft-in-reset, 2168
 disable bgp peer-group use-ip-router-alert, 2169
 disable bootp vlan, 1775
 disable bootprelay, 1776
 disable clear-flow, 1306

disable cli scripting, 495
 disable cli space-completion, 95
 disable cli-config-logging, 718
 disable clipaging, 96
 disable cna-testplug, 2466
 disable cpu-dos-protect, 1131
 disable cpu-monitoring, 275
 disable dhcp ports vlan, 1130
 disable dhcp vlan, 188
 disable diffserv examination port, 1029
 disable diffserv replacement port, 1030
 disable dot1p examination inner-tag port, 853
 disable dot1p examination ports, 1031
 disable dot1p replacement ports, 1032
 disable eaps, 1387
 disable edp ports, 372
 disable elrp-client, 2430
 disable elsm ports, 719
 disable elsm ports auto-restart, 720
 disable esrp, 1540
 disable flooding port, 903
 disable icmp address-mask, 1777
 disable icmp parameter-problem, 1778
 disable icmp port-unreachables, 1779
 disable icmp redirects, 1780
 disable icmp time-exceeded, 1781
 disable icmp timestamp, 1782
 disable icmp unreachable, 1783
 disable icmp userredirects, 1784
 disable idletimeout, 97
 disable igmp, 2246
 disable igmp snooping, 2247
 disable inline-power, 608
 disable inline-power legacy, 610
 disable inline-power legacy slot, 611
 disable inline-power ports, 612
 disable inline-power slot, 614
 disable iparp checking, 1785
 disable iparp gratuitous protect, 1132
 disable iparp refresh, 1786
 disable ipforwarding, 1787
 disable ipforwarding ipv6, 1872
 disable ipmcforwarding, 2249
 disable ip-option loose-source-route, 1789
 disable ip-option record-route, 1790
 disable ip-option record-timestamp, 1791
 disable ip-option strict-source-route, 1792
 disable ip-security anomaly-protection, 1133
 disable ip-security anomaly-protection icmp, 1138
 disable ip-security anomaly-protection ip, 1134
 disable ip-security anomaly-protection l4port, 1135
 disable ip-security anomaly-protection notify, 1139
 disable ip-security anomaly-protection tcp flags, 1136
 disable ip-security anomaly-protection tcp fragment, 1137
 disable ip-security arp gratuitous-protection vlan, 1140
 disable ip-security arp learning learn-from-arp vlan ports, 1141
 disable ip-security arp learning learn-from-dhcp vlan ports, 1143
 disable ip-security arp validation, 1144
 disable ip-security dhcp-snooping, 1145
 disable ip-security source-ip-lockdown ports, 1146
 disable irdp, 1794
 disable jumbo-frame ports, 373
 disable learning port, 375
 disable lldp ports, 580
 disable log debug-mode, 722, 2431
 disable log display, 723
 disable log target, 724
 disable loopback-mode vlan, 854
 disable mirroring, 376
 disable mld, 2296
 disable mld forward-mcrouter-only, 2297
 disable mld snooping, 2297
 disable mld snooping with-proxy, 2297
 disable mpls, 1643, 1656
 disable mpls ldp, 1646
 disable mpls rsvp-te summary-refresh, 1655
 disable msdp, 2337
 disable msdp data-encapsulation, 2338
 disable msdp export local-sa, 2339
 disable msdp peer, 2340
 disable msdp process-sa-request, 2341
 disable mvr, 2250
 disable netlogin, 1279
 disable netlogin dot1x guest-vlan ports, 1280
 disable netlogin logout-privilege, 1281
 disable netlogin ports, 1282
 disable netlogin session-refresh, 1283
 disable ospf, 2001
 disable ospf capability opaque-lsa, 2002
 disable ospf export, 2003
 disable ospf mpls-next-hop, 1657
 disable ospf originate-default, 2004
 disable ospf restart-helper-lsa-check, 2005
 disable ospf use-ip-router-alert, 2006
 disable ospfv3, 2062
 disable ospfv3 export, 2063
 disable peer-group, 2165
 disable pim, 2251

disable pim ssm vlan, 2252
 disable port, 377
 disable radius, 1148
 disable radius-accounting, 1149
 disable rip, 1902
 disable rip aggregation, 1903
 disable rip export, 1904
 disable rip originate-default, 1905
 disable rip poisonreverse, 1906
 disable rip splithorizon, 1907
 disable rip triggerupdate, 1908
 disable rip use-ip-router-alert, 1909
 disable ripng, 1940
 disable ripng export, 1941
 disable ripng originate-default, 1942
 disable ripng poisonreverse, 1943
 disable ripng splithorizon, 1944
 disable ripng triggerupdate, 1945
 disable rmon, 726
 disable router-discovery, 1874
 disable sflow, 727
 disable sflow ports, 728
 disable sharing, 378
 disable slot, 379
 disable smartredundancy, 381
 disable snmp access, 189
 disable snmp traps, 190
 disable snmp traps lldp, 581, 582
 disable snmp traps mpls, 1658, 1678
 disable snmp traps port-up-down ports, 382
 disable snmp-client, 191
 disable ssh2, 1150
 disable stacking, 1333
 disable stpd, 1464
 disable stpd auto-bind, 1465
 disable stpd ports, 1466
 disable stpd rapid-root-failover, 1467
 disable sys-health-check, 729
 disable syslog, 731
 disable tacacs, 1151
 disable tacacs-accounting, 1152
 disable tacacs-authorization, 1153
 disable telnet, 192
 disable udp-echo-server, 1795
 disable upm profile, 477
 disable use-ip router-alert, 1793
 disable vlan, 855
 disable vpls, 1583, 1659
 disable vpls service, 1660
 disable vrrp vrid, 1568
 disable watchdog, 193
 disable web http, 1154
 disable web https, 1155
 disable xml-mode, 276

download bootrom, 2365
 download image, 2368
 download ssl certificate, 1156
 download ssl privkey, 1158

E

edit policy, 927
 edit upm profile, 478
 eject memorycard, 2432
 ELSE, 496
 enable access-list permit to-cpu, 951
 enable access-list refresh blackhole, 952
 enable bgp, 1661, 2170
 enable bgp aggregation, 2171
 enable bgp always-compare-med, 2172
 enable bgp community format, 2173
 enable bgp export, 2174
 enable bgp neighbor, 2176
 enable bgp neighbor capability, 2177
 enable bgp neighbor remove-private-AS-numbers, 2179
 enable bgp neighbor soft-in-reset, 2180
 enable bgp neighbor use-ip-router-alert, 2181
 enable bgp peer-group, 2182
 enable bgp peer-group capability, 2183
 enable bgp peer-group remove-private-AS-number, 2184
 enable bgp peer-group soft-in-reset, 2185
 enable bgp peer-group use-ip-router-alert, 2186
 enable bootp vlan, 1796
 enable bootprelay, 1797
 enable clear-flow, 1307
 enable cli scripting, 497
 enable cli space-completion, 98
 enable cli-config-logging, 732
 enable clipaging, 99
 enable cna-testplug, 2467
 enable cpu-monitoring, 277
 enable dhcp ports vlan, 1160
 enable dhcp vlan, 194
 enable diffserv examination port, 1033
 enable diffserv replacement ports, 1034
 enable dos-protect, 1161
 enable dos-protect simulated, 1162
 enable dot1p examination inner-tag port, 856
 enable dot1p examination ports, 1036
 enable dot1p replacement ports, 1037
 enable eaps, 1389
 enable edp ports, 383
 enable elrp-client, 2433
 enable elsm ports, 733
 enable elsm ports auto-restart, 735
 enable esrp, 1541, 1545

enable flooding port, 906
 enable icmp address-mask, 1798
 enable icmp parameter-problem, 1799
 enable icmp port-unreachables, 1800
 enable icmp redirects, 1801
 enable icmp time-exceeded, 1802
 enable icmp timestamp, 1803
 enable icmp unreachable, 1804
 enable icmp userredirects, 1805
 enable idletimeout, 100
 enable igmp, 2253
 enable igmp snooping, 2254
 enable igmp snooping with-proxy, 2256
 enable inline-power, 615
 enable inline-power legacy, 617
 enable inline-power legacy slot, 618
 enable inline-power ports, 620
 enable inline-power slot, 622
 enable iparp checking, 1806
 enable iparp gratuitous protect, 1163
 enable iparp refresh, 1807
 enable ipforwarding, 1808
 enable ipforwarding ipv6, 1875
 enable ipmcforwarding, 2257
 enable ip-option loose-source-route, 1810
 enable ip-option record-route, 1811
 enable ip-option record-timestamp, 1812
 enable ip-option router-alert, 1814
 enable ip-option strict-source-route, 1813
 enable iproute sharing, 1816
 enable ip-security anomaly-protection, 1164
 enable ip-security anomaly-protection icmp, 1165
 enable ip-security anomaly-protection ip, 1166
 enable ip-security anomaly-protection l4port, 1167
 enable ip-security anomaly-protection notify, 1168
 enable ip-security anomaly-protection tcp flags, 1169
 enable ip-security anomaly-protection tcp fragment, 1170
 enable ip-security arp gratuitous-protection vlan, 1171
 enable ip-security arp learning learn-from-arp vlan ports, 1173
 enable ip-security arp learning learn-from-dhcp vlan ports, 1175
 enable ip-security arp validation vlan violation-action, 1177
 enable ip-security dhcp-snooping vlan ports violation-action, 1179
 enable ip-security source-ip-lockdown ports, 1181
 enable irdp, 1817
 enable jumbo-frame ports, 385
 enable learning port, 387
 enable license file, 104
 enable license software, 101
 enable lldp ports, 583
 enable log debug-mode, 737, 2434
 enable log display, 738
 enable log target, 740
 enable loopback-mode vlan, 857
 enable mac-lockdown-timeout ports, 1183
 enable mirroring to port, 388
 enable mld, 2298
 enable mld snooping, 2299
 enable mld snooping forward-mcrouter-only, 2299
 enable mld snooping with-proxy, 2301
 enable mpls, 1644, 1663, 1664
 enable mpls ldp bgp-routes, 1667
 enable mpls php vlan, 1649, 1669
 enable mpls protocol ldp, 1648, 1650, 1666, 1668, 1670
 enable mpls rsvp-te lsp, 1674
 enable mpls rsvp-te summary-refresh, 1675
 enable msdp, 2343
 enable msdp data-encapsulation, 2344
 enable msdp export local-sa, 2345
 enable msdp peer, 2347
 enable msdp process-sa-request, 2348
 enable mvr, 2258
 enable netlogin, 1284
 enable netlogin dot1x guest-vlan ports, 1285
 enable netlogin logout-privilege, 1287
 enable netlogin ports, 1288
 enable netlogin session-refresh, 1289
 enable ospf, 2007, 2064
 enable ospf capability opaque-lsa, 2008
 enable ospf export, 2009
 enable ospf mpls-next-hop, 1677
 enable ospf originate-default, 2011
 enable ospf restart-helper-lsa-check, 2012
 enable ospf use-ip-router-alert, 2013
 enable ospfv3 export, 2065
 enable pim, 2259
 enable pim ssm vlan, 2260
 enable port, 392
 enable radius, 1184
 enable radius-accounting, 1186
 enable rip, 1910
 enable rip aggregation, 1911
 enable rip export, 1912
 enable rip originate-default cost, 1914
 enable rip poisonreverse, 1915
 enable rip splithorizon, 1916
 enable rip triggerupdate, 1917

enable rip use-ip-router-alert, 1918
 enable ripng, 1946
 enable ripng export, 1947
 enable ripng originate-default, 1949
 enable ripng poisonreverse, 1950
 enable ripng splithorizon, 1951
 enable ripng triggerupdate, 1952
 enable rmon, 742
 enable router-discovery, 1877
 enable sflow, 744
 enable sflow ports, 746
 enable sharing grouping, 393
 enable slot, 398
 enable smartredundancy, 400
 enable snmp access, 195
 enable snmp traps, 197
 enable snmp traps lldp, 584
 enable snmp traps lldp-med, 585
 enable snmp traps port-up-down ports, 401
 enable snmp-client, 198
 enable ssh2, 1188
 enable stacking, 1335
 enable stpd, 1468
 enable stpd auto-bind, 1469
 enable stpd ports, 1472
 enable stpd rapid-root-failover, 1473
 enable sys-health-check, 748
 enable syslog, 750
 enable tacacs, 1190
 enable tacacs accounting, 1191
 enable tacacs-authorization, 1192
 enable telnet, 199
 enable udp-echo-server, 1818
 enable upm profile, 479
 enable use-ip-router-alert, 1814
 enable vlan, 858
 enable vpls, 1679
 enable vpls service, 1680
 enable vrrp vrid, 1569
 enable watchdog, 201
 enable web http, 1193
 enable web https, 1194
 enable xml-mode, 278
 ENDIF, 498
 ENDWHILE, 499
 exit, 202

H

history, 106

I

install bootrom, 2380

install firmware, 2383
 install image, 2387

L

load script, 2391
 logout, 203
 ls, 279

M

mv, 282

N

nslookup, 2435

P

ping, 107
 ping mac port, 533
 ping mpls lsp, 1681

Q

quit, 204

R

reboot, 109
 refresh policy, 929
 reset inline-power ports, 624
 restart ports, 402
 restart process, 285
 rm, 288
 rtlookup, 1819, 1878
 run elrp, 2442
 run msm-failover, 403, 405
 run update, 2393

S

save configuration, 2394
 save debug tracefiles, 2444
 scp2, 1196
 set var, 504
 show access-list, 953
 show access-list configuration, 955
 show access-list counter, 956
 show access-list dynamic, 958
 show access-list dynamic counter, 959
 show access-list dynamic rule, 960
 show access-list meter, 965, 1039
 show access-list usage acl-mask, 966
 show access-list usage acl-range, 967

show access-list usage acl-rule, 969
 show access-list usage acl-slice, 971
 show accounts, 111
 show accounts password-policy, 113
 show banner, 115
 show banner netlogin, 1290
 show bgp, 2187
 show bgp memory, 2197
 show bgp neighbor, 2188
 show bgp neighbor accepted-routes, 2191
 show bgp neighbor flap-statistics, 2189
 show bgp neighbor received-routes, 2191
 show bgp neighbor rejected-routes, 2191
 show bgp neighbor routes, 2191
 show bgp neighbor suppressed-routes, 2191
 show bgp neighbor transmitted-routes, 2191
 show bgp peer-group, 2193
 show bgp routes, 2194
 show bgp routes summary, 2196
 show bootprelay, 1820
 show bvlan, 859
 show cfm, 535
 show cfm detail, 537
 show cfm statistics, 539
 show checkpoint-data, 206
 show clear-flow, 1308
 show clear-flow acl-modified, 1309
 show clear-flow rule, 1310
 show clear-flow rule-all, 1312
 show clear-flow rule-triggered, 1314
 show cna-testplug, 2468
 show configuration, 2397
 show debug, 2445
 show dhcp-client state, 205
 show dhcp-server, 1199
 show diagnostics, 2447
 show diffserv, 1040
 show dns-client, 116
 show dos-protect, 1200, 1207, 1208, 1209, 1211
 show dot1p, 1042
 show eaps, 1391
 show eaps counters, 1395
 show eaps counters shared-port, 1400
 show eaps shared-port, 1405
 show edp, 406
 show elrp, 2451
 show elsm, 751
 show elsm ports, 754
 show esrp, 1542
 show esrp counters, 1547
 show fans, 759
 show fdb, 908
 show fdb vpls, 1687
 show forwarding configuration, 2453
 show heartbeat process, 295
 show igmp, 2261
 show igmp group, 2263
 show igmp snooping, 2265
 show igmp snooping vlan filter, 2267
 show igmp snooping vlan static, 2268
 show inline-power, 625
 show inline-power configuration ports, 627
 show inline-power info ports, 629
 show inline-power slot, 632
 show inline-power stats, 634
 show inline-power stats ports, 635
 show inline-power stats slot, 637
 show iparp, 1821
 show iparp proxy, 1824
 show ipconfig, 1825
 show ipconfig ipv6, 1879
 show iproute, 1826
 show iproute ipv6, 1880
 show iproute ipv6 origin, 1881
 show iproute origin, 1828
 show ip-security anomaly-protection notify cache ports, 1202
 show ip-security arp gratuitous-protection, 1203
 show ip-security arp learning vlan, 1204
 show ip-security arp validation, 1206
 show ip-security dhcp-snooping entries vlan, 1207
 show ip-security dhcp-snooping violations vlan, 1208
 show ip-security dhcp-snooping vlan, 1209
 show ip-security gratuitous-protection vlan, 1171
 show ip-security source-ip-lockdown, 1211
 show ipstats, 1829, 1882
 show L2stats, 2271
 show l2stats, 2271
 show lacp, 408
 show lacp counters, 410
 show lacp lag, 412
 show lacp member-port, 417
 show licenses, 117
 show lldp, 586
 show lldp neighbors, 589
 show lldp statistics, 592
 show log, 763
 show log components, 767
 show log configuration, 771
 show log configuration filter, 773
 show log configuration target, 775
 show log counters, 778
 show log events, 780
 show mac-binding, 861
 show mac-lockdown-timeout fdb ports, 1212

show mac-lockdown-timeout ports, 1214
show management, 208
show memory, 298
show memory process, 303
show memorycard, 2400
show meter, 1044
show mirroring, 420
show mld, 2302
show mld group, 2304
show mld snooping, 2305
show mld snooping vlan static, 2307
show mpls, 1689, 1706
show mpls exp replacement, 1692
show mpls label, 1693
show mpls ldp, 1704
show mpls ldp interface, 1700, 1710
show mpls ldp lsp, 1702
show mpls rsvp-te bandwidth, 1708
show msdp, 2350
show msdp memory, 2351
show msdp mesh-group, 2353
show msdp peer, 2355
show msdp sa-cache, 2357
show mvr, 2272
show mvr cache, 2273
show neighbor-discovery cache ipv6, 1883
show netlogin, 1291
show netlogin guest-vlan, 1297
show netlogin local-users, 1298
show netlogin mac-list, 1300
show netlogin ports, 1291
show node, 211
show odometers, 213
show ospf, 2014
show ospf area, 2015
show ospf area detail, 2016
show ospf ase-summary, 2017
show ospf interfaces, 2019
show ospf interfaces detail, 2018
show ospf memory, 2022
show ospf neighbor, 2023
show ospf virtual-link, 2024
show ospfv3, 2067
show ospfv3 area, 2069
show ospfv3 interfaces, 2071
show ospfv3 memory, 2077
show ospfv3 neighbor, 2078
show ospfv3 virtual-link, 2079
show pim, 2274
show pim cache, 2277
show pim rp-set, 2274
show policy, 931
show ports anomaly, 422
show ports collisions, 424
show ports configuration, 426
show ports information, 429
show ports packet, 435
show ports qosmonitor, 1047
show ports redundant, 437
show ports rxerrors, 782
show ports sharing, 438
show ports statistics, 785
show ports txerrors, 788
show ports utilization, 440
show ports vlan statistics, 791
show ports vlan stats, 791
show ports wan-phy configuration, 445
show ports wan-phy errors, 447
show ports wan-phy events, 449
show ports wan-phy overhead, 451
show power, 216, 1337
show power budget, 220
show power controller, 223
show powersupplies, 216
show process, 306
show protocol, 863
show qosprofile, 1049
show radius, 1216
show radius-accounting, 1219
show rip, 1919
show rip interface, 1920
show rip interface vlan, 1922
show rip memory, 1923
show rip routes, 1924
show ripng, 1953
show ripng interface, 1954
show ripng routes, 1956
show rmon memory, 793
show router-discovery, 1885
show session, 225
show sflow configuration, 796
show sflow statistics, 798
show slot, 453
show snmpv3 access, 227
show snmpv3 context, 230
show snmpv3 counters, 231
show snmpv3 engine-info, 233
show snmpv3 extreme target-addr-extr, 246
show snmpv3 filter, 234
show snmpv3 filter-profile, 236
show snmpv3 group, 238
show snmpv3 mib-view, 241
show snmpv3 notify, 243
show snmpv3 target-addr, 244
show snmpv3 target-params, 247
show snmpv3 user, 249
show snmp-client, 251
show sshd2 user-key, 1221

show ssl, 1222
 show stacking, 1339
 show stacking configuration, 1342
 show stacking detail, 1344
 show stacking stack-ports, 1346
 show stpd, 1474
 show stpd ports, 1478
 show svlan, 864
 show switch, 119
 show tacacs, 1224
 show tacacs-accounting, 1226
 show tech, 2454
 show temperature, 800
 show traffic mode, 1052
 show traffic queue, 1053
 show traffic queue statistics, 1056
 show traffic queue utilization, 1059
 show tunnel, 1886
 show udp-profile, 1832
 show version, 803
 show virtual-router, 918
 show vlan, 866
 show vlan dhcp-address-allocation, 1228
 show vlan dhcp-config, 1229
 show vlan eaps, 1410
 show vlan security, 1230
 show vlan statistics, 808
 show vlan stats, 808
 show vlan stpd, 1481
 show vman, 893
 show vpls, 1729
 show vpls peer, 1731
 show vr, 918
 show vrrp, 1570
 show vrrp vir, 1575
 show vrrp virtual-router, 1573
 show vrrp vlan, 1577
 ssh2, 1231
 start process, 313
 synchronize, 2401
 synchronize stacking, 1348

T

telnet, 253
 telnet msm, 255
 telnet slot, 257
 terminate process, 315
 tftp, 259
 tftp get, 263
 tftp put, 266
 top, 2456
 traceroute, 123
 traceroute mac port, 541

traceroute mpls lsp, 1684

U

unconfigure access-list, 973
 unconfigure bootprelay dhcp-agent information check, 1833
 unconfigure bootprelay dhcp-agent information option, 1834
 unconfigure bootprelay dhcp-agent information policy, 1835
 unconfigure diffserv, 1061
 unconfigure eaps port, 1414
 unconfigure eaps shared-port link-id, 1412
 unconfigure eaps shared-port mode, 1413
 unconfigure elrp-client, 2457
 unconfigure icmp, 1836
 unconfigure igmp, 2279
 unconfigure inline-power budget slot, 639
 unconfigure inline-power disconnect-precedence, 640
 unconfigure inline-power operator-limit ports, 641
 unconfigure inline-power priority ports, 642
 unconfigure inline-power usage-threshold, 643
 unconfigure iparp, 1837
 unconfigure irdp, 1838
 unconfigure lldp, 594
 unconfigure log filter, 809
 unconfigure log target format, 810
 unconfigure mld, 2308
 unconfigure mpls vlan, 1735, 1738
 unconfigure msdp sa-cache-server, 2359
 unconfigure mstp region, 1483
 unconfigure netlogin authentication database-order, 1301
 unconfigure netlogin dot1x guest-vlan, 1302
 unconfigure netlogin vlan, 1303
 unconfigure ospf, 2025
 unconfigure ospfv3, 2080
 unconfigure pim, 2280
 unconfigure pim border, 2360
 unconfigure pim ssm range, 2281
 unconfigure ports aggregate-bandwidth, 1062
 unconfigure ports display-string, 462
 unconfigure ports monitor vlan, 812
 unconfigure ports redundant, 463
 unconfigure ports wan-phy, 464
 unconfigure qosprofile, 1063
 unconfigure radius, 1234
 unconfigure radius-accounting, 1236
 unconfigure rip, 1925
 unconfigure ripng, 1958
 unconfigure sflow, 813
 unconfigure sflow agent, 814

- unconfigure sflow collector ipaddress, 815
- unconfigure sflow ports, 816
- unconfigure slot, 465
- unconfigure stacking alternate-ip-address, 1352
- unconfigure stacking license-level, 1353
- unconfigure stpd, 1484
- unconfigure stpd ports link-type, 1485
- unconfigure switch, 2406
- unconfigure tacacs, 1238
- unconfigure tacacs-accounting, 1239
- unconfigure tunnel, 1887
- unconfigure udp-profile, 1839
- unconfigure upm event, 488
- unconfigure upm timer, 489
- unconfigure vlan dhcp, 1241
- unconfigure vlan dhcp-address-range, 1242
- unconfigure vlan dhcp-options, 1243
- unconfigure vlan ipaddress, 870
- unconfigure vman rate-limit flood-traffic
designated-port, 1066
- unconfigure vpls, 1739
- uninstall image, 2408
- upload configuration, 2410
- upload debug, 2458
- upload debug-info memorycard, 2444
- upload log, 817
- use configuration, 2415
- use image, 2417

V

- virtual-router, 920

W

- WHILE ... DO, 507