SOFTWARE MANUAL

Mellanox Onyx User Manual

Rev 5.7
Software Version 3.7.11xx

NOTE:
THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

Table of Contents

List of Tables

List of Figures

Document Revision History

Rev 5.7 – December 31, 2018

Updated the following commands/sections:

- “system profile” on page 540
- “show what-just-happened” on page 705
- Section 5.8.1, “MLAG Keepalive and Failover,” on page 854
- “link state tracking group” on page 884
- “link state tracking vlan” on page 885
- “switchport access” on page 801
- “Signal Degradation Monitoring” on page 377
- Section 6.1.2.2, “ECMP Consistent Hashing,” on page 1216
- “ip load-sharing” on page 1275
- “show ip load-sharing” on page 1286
- “show ip route” on page 1276
- Section 5.1.1.1, “Changing the Module Type to a Split Mode,” on page 707
- Table 56, “Port Splitting Options”
- “bfd interval” on page 1461
- “show ip route static” on page 1464
- “set community” on page 1484
- “magp” on page 1575
- “vrrp” on page 1562
- Section 5.20.1, “Configuring UDK,” on page 1147
- Section 4.24, “What Just Happened (WJH),” on page 702
- Section 5.1.2, “56GbE Link Speed,” on page 708
- “show interfaces ethernet” on page 730
- “ip pim multipath next-hop” on page 1520
- “show ip pim protocol” on page 1523
- “aaa authentication login” on page 437
- “stats sample interval” on page 518
- “stats export” on page 516
- “ip route bfd” on page 1463
- “ip igmp last-member-query-response-time” on page 1545
- “ip igmp snooping (config)” on page 1029
- Section 4.24, “What Just Happened (WJH),” on page 702
- “what-just-happened” on page 703

- “show what-just-happened” on page 705
- “show ip pim rp-hash” on page 1530
- “show ssh client source-interface” on page 204
- Section 6.3, “OSPF,” on page 1331
- Section 6.3.5.2,“Config Router,” on page 1338
- Section 6.3.5.3, “Interface,” on page 1351
- “stats sample <sample-id> clear” on page 514
- “stats sample <sample-id> enable” on page 515
- “show stats sample” on page 523
- “show stats sample data” on page 524
- Section 6.4.8.3, “Show,” on page 1435

Added:

- “ip pim multipath rp” on page 1521
- “ipv6 dhcp client enable” on page 159
- “ipv6 dhcp client renew” on page 160
- “stats sample max-entries” on page 517
- “show stats sample data” on page 524

Rev 5.6 – December 13, 2018

Added:

- Section 4.2, “Management Source IP Address,” on page 188

Rev 5.5 – December 04, 2018

Added:

- the command “clear ptp interface port-channel counters” on page 248
- the command “clear ptp VRF counters” on page 249
- the command “interface port-channel” on page 250
- the command “ptp vrf” on page 251
- the command “show ptp interface port-channel” on page 259
- the command “show ptp vrf” on page 253
- the command “show ptp vrf counters” on page 255
- the command “show ptp interface port-channel counters” on page 260
- the command “email autosupport mailhub” on page 394
- the command “email autosupport recipient” on page 395
- the command “show email” on page 405
- the command “snmp-server cache enable” on page 592

- Section 4.24, “What Just Happened (WJH),” on page 702
- Section 5.9, “Link State Tracking,” on page 881
- Action “Policer” added to Table 53, “Summary of ACL Capability,” on page 892

Updated:

- Appendix 6.9,“MAGP,” on page 1572
- Section 4.1.7.7, “IP Diagnostic Tools,” on page 184
- Section 4.4.4, “Configuring PTP,” on page 229
- the command “show ptp forced-master” on page 265
- the command “show ptp” on page 252
- Section 4.11.1, “Supported Events,” on page 380
- the command “aaa authorization” on page 446
- the command “show aaa” on page 448
- Section 4.14.1, “System File Encryption,” on page 475
- Table 29, “System Health Monitor Alerts Scenarios,” on page 528
- the command “system profile” on page 540
- the command “show memory” on page 549
- the command “Configuring an SNMPv3 User” on page 570
- the command “snmp-server user” on page 609
- the command “show snmp auto-refresh” on page 612
- the command “show puppet-agent” on page 636
- the command “show virtual-machine interface” on page 662
- Section 4.22, “Resource Scale,” on page 678
- Section 5.1.2, “56GbE Link Speed,” on page 708
- the command “fec-override” on page 716
- the command “show interfaces ethernet rates” on page 733
- the command “show interfaces port-channel” on page 767
- Section 5.6.2, “Port Type,” on page 805
- Section 5.6.4, “BPDU Guard,” on page 806
- Section 5.6.5, “Loop Guard,” on page 806
- the command “spanning-tree mst root” on page 826
- Section 5.9.1, “Configuring Link State Tracking,” on page 881
- the command “link state tracking group” on page 884
- the command “link state tracking vlan” on page 885
- the command “deny/permit (MAC ACL rule)” on page 903
- the command “deny/permit (IPv4 ACL rule)” on page 905
- the command “deny/permit (IPv4 TCP ACL rule)” on page 907

- the command “deny/permit (IPv4 TCP-UDP/UDP ACL rule)” on page 910
- the command “deny/permit (IPv4 ICMP ACL rule)” on page 912
- the command “deny/permit (IPv6 ACL rule)” on page 914
- the command “deny/permit (IPv6 TCP ACL rule)” on page 916
- the command “deny/permit (IPv6 TCP-UDP/UDP ACL rule)” on page 918
- the command “deny/permit (IPv6 ICMPv6 ACL rule)” on page 920
- the command “deny/permit (MAC UDK ACL rule)” on page 922
- the command “deny/permit (IPv4 UDK ACL rule)” on page 924
- the command “deny/permit (IPv4 TCP UDK ACL rule)” on page 926
- the command “deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)” on page 929
- the command “deny/permit (IPv4 ICMP UDK ACL rule)” on page 932
- the command “show access-lists action” on page 946
- Section 5.13.1, “Configuring VXLAN,” on page 994
- Section 5.14.3, “IGMP Snooping Querier,” on page 1026
- the command “igmp snooping querier query-interval” on page 1035
- the command “Trust Levels” on page 1056
- the command “qos default switch-priority” on page 1064
- the command “storm-control” on page 1142
- Section 6.1.1.6, “Configuring a Router Port Interface,” on page 1212
- the command “show ip” on page 1251
- the command “show ip interface port-channel” on page 1256
- the command “show ip interface vrf” on page 1259
- Section 6.3.3, “Configuring OSPF,” on page 1332
- Section 6.4.5, “Configuring BGP,” on page 1375
- the command “show {ip | ipv6} bgp” on page 1435

Removed Appendix “Show Commands Not Supported by JSON”

Rev 5.4 – November 5, 2018

No changes made since last revision.

Rev 5.3 – August 30, 2018

Added:

- the command “web proxy auth authtype” on page 136
- the command “web proxy auth basic” on page 137
- the command “web proxy auth host” on page 138

Updated:

- the command "{ip | ipv6} route" on page 171
- the command "image install" on page 293
- the command "image options" on page 295
- Section 4.13.2, "Authentication, Authorization and Accounting (AAA)," on page 427
- the command "aaa authorization" on page 446
- the command "show virtual-machine install" on page 661
- the command "show telemetry" on page 420
- the command "start" on page 692
- the command "speed" on page 723
- the command "show mac access-lists summary" on page 944
- the command "dcb priority-flow-control mode" on page 1051
- the command "show buffers details" on page 1134
- Section 6.4.8, "BGP Commands," on page 1378
- the command "show ip bgp address-family" on page 1437
- the command "show ip bgp neighbors" on page 1443
- the command "show ip bgp neighbors received" on page 1445
- the command "show ip bgp neighbors received detail" on page 1446
- the command "show ip bgp peer-group" on page 1448
- the command "show ip bgp update-group" on page 1451
- the command "vrrp" on page 1562
- the command "ip virtual-router address" on page 1577

Rev 5.2 – August 1, 2018

Added:

- Section 4.4.4, "Configuring PTP," on page 229
- Section 4.4.5, "Securing PTP Infrastructure," on page 231
- the command "ptp amt" on page 235
- the command "ptp enable forced-master" on page 241
- the command "clear ptp amt log" on page 245
- the command "clear ptp forced-master log" on page 246
- the command "show ptp amt" on page 258
- the command "show ptp amt log" on page 262
- the command "show ptp forced-master" on page 265
- the command "show ptp forced-master log" on page 266
- Section 4.5.1, "Important Pre-OS Upgrade Notes," on page 274
- Section 5.12.4, "Configuring Flows Using CLI Commands," on page 962

- the command “[neighbor] next-hop-unchanged” on page 1409
- the command “neighbor soft-reconfiguration inbound” on page 1423
- the command “ip pim multipath next-hop” on page 1520
- the command “ip igmp immediate-leave” on page 1544

Removed SwitchX content from the document

Updated:

- Table 4, “Ethernet Features,” on page 41
- the command “hostname” on page 165
- the command “show ptp” on page 252
- the command “show ptp” on page 267
- the command “show ptp interface vlan” on page 269
- the command “telemetry sampling tc” on page 410
- the command “show interfaces ethernet” on page 730
- the command “lldp tlv-select” on page 782
- the command “openflow add-flows” on page 968
- the command “openflow add-group” on page 971
- the command “openflow mod-group” on page 973
- the command “openflow add-meter” on page 974
- the command “openflow mod-meter” on page 976
- the command “openflow re-apply flows” on page 977
- the command “openflow re-apply groups” on page 978
- the command “openflow re-apply meters” on page 979
- the command “show ip route” on page 1276
- the command “address-family” on page 1383
- the command “neighbor activate” on page 1395
- the note section of command “neighbor advertisement-interval” on page 1396
- the command “neighbor route-map” on page 1415
- the command “neighbor route-reflector-client” on page 1418
- the command “neighbor send-community” on page 1419
- the command “show ip bgp address-family” on page 1437
- the command “show ip bgp evpn” on page 1440
- the command “show ip bgp evpn summary” on page 1442
- the command “show ip bgp neighbors” on page 1443
- the command “show ip bgp peer-group” on page 1448
- the command “show ip bgp vrf summary” on page 1452
- Section 6.7.4.3, “Configuring PIM,” on page 1505

- the command “ip pim bsr-candidate” on page 1510
- the command “vrrp” on page 1562
- the command “show ip igmp interface” on page 1554
- the command “ip virtual-router address” on page 1577
- the command “show magp” on page 1579
- the command “show magp interface vlan” on page 1580

Rev 5.1 – June 28, 2018

Added:

- Section 3.3.7, “IP Route,” on page 84
- the command “web https ssl renegotiation enable” on page 134
- the command “web https ssl secure-cookie enable” on page 135
- the command “show interfaces mgmt0” on page 161
- the command “show interfaces mgmt0 brief” on page 163
- the command “show interfaces mgmt0 configured” on page 164
- the command “ptp message-format” on page 242
- the command “clear ptp interface vlan ethernet counters” on page 247
- the command “show ptp interface vlan” on page 269
- the command “show ptp interface vlan ethernet” on page 270
- the command “show ptp interface vlan counters” on page 271
- the command “show ptp interface vlan ethernet counters” on page 273
- the command “ldap hostname-check enable” on page 461
- the command “show ldap crl” on page 472
- Table 36, “SNMP MELLANOX-POWER-CYCLE Traps,” on page 569
- Section 4.18.1.5, “Resetting SNMPv3 Engine ID,” on page 569
- the command “show snmp” on page 611
- the command “show snmp engineID” on page 613
- the command “show docker containers” on page 696
- the command “show docker stats” on page 701
- Section 5.8.6, “Interoperability with MLAG,” on page 855
- the command “nve fdb flood load-balance” on page 1006
- the command “nve fdb learning remote” on page 1007
- the command “nve vlan bridge” on page 1009
- the command “clear mac-address-table nve” on page 1012
- Section 6.1.3, “ARP Neighbor Discovery Responder,” on page 1219
- the command “show interfaces” on page 1245

- the command “show interfaces vlan” on page 1246
- the command “show ip interface” on page 1247
- the command “show ip interface brief” on page 1249
- the command “show interface configured” on page 1250
- the command “show ip” on page 1251
- the command “show ip interface mgmt0” on page 1254
- the command “show ipv6 interface” on page 1261
- the command “show ipv6 interface brief” on page 1263
- the command “show ipv6” on page 1264
- the command “show ipv6 interface loopback” on page 1265
- the command “show ipv6 interface port-channel” on page 1266
- the command “show ipv6 interface vlan” on page 1267
- the command “show ipv6 interface vrf” on page 1268
- the command “ip arp responder” on page 1288
- the command “default-information originate” on page 1349
- the command “clear ip pim counters” on page 1522
- the command “show ip pim interface brief” on page 1527
- the command “show ip igmp groups” on page 1553
- the command “show ip igmp interface” on page 1554
- the command “show ip igmp interface brief” on page 1556
- the command “use-secondary-ip” on page 1588
- the command “ipv6 dhcp relay instance use-secondary-ip” on page 1603
- Appendix B, “Show Commands Not Supported by JSON” on page 1943

Updated:

- the command “show web” on page 139
- the command “ptp enable” on page 240
- the command “ptp domain” on page 239
- the command “ptp priority” on page 243
- the command “ptp announce interval” on page 236
- the command “ptp announce timeout” on page 237
- the command “ptp delay-req interval” on page 238
- the command “ptp sync interval” on page 244
- the command “show ptp” on page 252
- Table 26, “Supported Event Notifications and MIB Mapping,” on page 380
- the command “show files stats telemetry” on page 426
- the command “ldap ssl” on page 466

- the command “show ldap” on page 471
- Section 4.18.2.7.1, “To Execute a JSON Request,” on page 588
- Section 4.18.2.7.2, “To Query an Asynchronous JSON Request,” on page 589
- the command “save” on page 690
- Section 4.23, “Linux Dockers,” on page 680
- the command “commit” on page 683
- the command “start” on page 692
- the command “show interfaces ethernet” on page 730
- the command “show lacp interfaces ethernet” on page 762
- the command “show interfaces port-channel compatibility-parameters” on page 771
- the command “show interfaces port-channel” on page 767
- the command “lldp tlv-select” on page 782
- Section 5.8.5, “Upgrading MLAG Pair,” on page 855
- the command “show mlag-vip” on page 875
- the command “show interfaces mlag-port-channel” on page 876
- the command “show access-lists log config” on page 948
- the command “openflow add-flows” on page 968
- the command “show ip igmp snooping” on page 1037
- the command “show ip igmp snooping groups” on page 1038
- the command “show qos” on page 1070
- the command “show qos” on page 1072
- the command “show qos interface port-channel” on page 1076
- the command “show qos interface rewrite-mapping” on page 1080
- the command “show ip routing” on page 1230
- the command “show ip interface vrf” on page 1259
- the command “show interfaces loopback” on page 1273
- the command “show ip route” on page 1276
- the command “show ip route vrf” on page 1278
- the command “show ip route failed” on page 1281
- the command “show ip route static” on page 1282
- the command “show ip route static multicast-override” on page 1283
- the command “show ipv6 interfaces brief” on page 1327
- the command “show ipv6 route” on page 1330
- the command “neighbor update-source” on page 1426
- Section 6.7.4.3, “Configuring PIM,” on page 1505
- the command “ip pim bsr-candidate” on page 1510

- the command “ip pim register-source” on page 1512
- the command “show ip pim protocol” on page 1523
- the command “show ip pim bsr” on page 1524
- the command “show ip pim interface” on page 1525
- the command “show ip pim neighbor” on page 1528
- the command “show ip pim rp” on page 1529
- the command “show ip pim rp-candidate” on page 1531
- the command “show ip pim ssm range” on page 1532
- the command “show ip pim upstream joins” on page 1533
- the configuration mode for the command “clear ip mroute” on page 1537
- the command “show ip mroute” on page 1538
- the command “show ip mroute summary” on page 1542
- the configuration mode for the command “clear ip igmp groups” on page 1552
- the command “show ip igmp groups” on page 1553
- the command “show ip igmp interface” on page 1554
- the command “show ip igmp interface brief” on page 1556
- the command “address” on page 1563
- the command “show vrrp detail” on page 1570
- the command “show ip dhcp relay” on page 1593
- the command “show ip dhcp relay counters” on page 1594
- the command “show ipv6 dhcp relay” on page 1605
- the command “show ipv6 dhcp relay counters” on page 1606
- Appendix A, “Enhancing System Security According to NIST SP 800-131A” on page 1607

Rev 5.0 – March 29, 2018

Added:

- the command “snmp-server engineID reset” on page 597
- the command “show interfaces counters discard” on page 729
- the command “show interfaces ethernet transceiver brief” on page 736
- the command “clear ip routing counters” on page 1229
- the command “show ip routing counters” on page 1231
- the command “ip pim sg-expiry-timer” on page 1508
- the command “ip pim register-source” on page 1512
- the command “clear ip mroute” on page 1537

Updated:

- Table 14, “WebUI Setup Submenus,” on page 80

- Step 5 in Section 4.18.1.7, “Configuring an SNMP Notification,” on page 571
- Table 51, “LR4/ER4 Switch and Port Support,” on page 711
- the command “flowcontrol” on page 717
- the command “show interfaces switchport” on page 804
- Step 1 in the procedure “To verify MLAG configuration:” on page 859.
- the command “show mlag” on page 874
- the command “show mlag statistics” on page 880
- the command “dcb priority-flow-control mode” on page 1051
- Section 5.17.2, “Lossless Traffic,” on page 1106
- the command “ip pim rp-candidate” on page 1513
- the command “show ip pim rp-hash” on page 1530
- the command “ip pim ssm range” on page 1519
- the command “show ip pim protocol” on page 1523
- the command “show ip pim bsr” on page 1524
- the command “show ip pim interface” on page 1525
- the command “show ip pim rp” on page 1529

Rev 4.90 – March 4, 2018

Added:

- Step 1 in section Section 2.5.1, “Installing Onyx License (CLI),” on page 60
- the command “show banner” on page 106
- the command “show ssh server host-keys” on page 122
- the command “show web” on page 139
- the command “image default-chip-fw” on page 289
- the command “logging event enable” on page 338
- the command “logging event error-threshold” on page 339
- the command “logging event interval” on page 340
- the command “logging event rate-limit” on page 341
- the command “show logging events” on page 357
- the command “show logging events source-counters” on page 359
- Section 4.17.7, “Viewing Active Events,” on page 533
- the command “clear system hardware events” on page 536
- the command “system profile” on page 540
- the command “show system hardware events” on page 556
- the command “show system profile detailed” on page 559
- the command “show ip filter” on page 672

- the command “show ip filter all” on page 673
- the command “show ip filter configured” on page 674
- the command “show ipv6 filter” on page 675
- the command “show ipv6 filter all” on page 676
- the command “show ipv6 filter configured” on page 677
- Section 5.12.2.1.2, “Non-standard Matches,” on page 958
- Section 5.15.2, “PFC Watchdog,” on page 1048
- the command “pfc-wd” on page 1052
- the command “show qos ip rewrite” on page 1088
- Section 5.17.3.8, “Exceptions to Legal Shared Buffer Configuration,” on page 1112
- the command “show ip route failed” on page 1281
- the command “show ip route static multicast-override” on page 1283
- the command “show ip bgp vrf summary” on page 1452
- the command “show ip pim ssm range” on page 1532
- Section 6.10.2, “Upstream and Downstream Interfaces,” on page 1581

Updated:

- Section 3.3.1, “Setup Menu,” on page 80
- Section 3.3.4, “Ports Menu,” on page 82
- Section 3.3.6, “ETH Mgmt,” on page 84
- the command “show ssh server” on page 121
- the command “show clock” on page 222
- the command “show ntp” on page 223
- the command “show debug ethernet” on page 372
- the command “show radius” on page 452
- the command “show tacacs” on page 456
- the command “show snmp auto-refresh” on page 612
- the command “speed” on page 723
- the command “show interfaces ethernet” on page 730
- the command “show interfaces ethernet transceiver diagnostics” on page 739
- the command “show lacp counters” on page 761
- the command “show interfaces port-channel compatibility-parameters” on page 771
- the command “switchport mode” on page 799
- the command “switchport access” on page 801
- the command “show spanning-tree” on page 833
- the command “show spanning-tree mst” on page 837
- the command “show interfaces mlag-port-channel” on page 876

- the command “show interfaces mlag-port-channel summary” on page 879
- the command “deny/permit (MAC ACL rule)” on page 903
- the command “deny/permit (IPv4 ACL rule)” on page 905
- the command “deny/permit (IPv4 TCP ACL rule)” on page 907
- the command “deny/permit (IPv4 TCP-UDP/UDP ACL rule)” on page 910
- the command “deny/permit (IPv4 ICMP ACL rule)” on page 912
- the command “deny/permit (IPv6 ACL rule)” on page 914
- the command “deny/permit (IPv6 TCP ACL rule)” on page 916
- the command “deny/permit (IPv6 TCP-UDP/UDP ACL rule)” on page 918
- the command “deny/permit (IPv6 ICMPv6 ACL rule)” on page 920
- the command “deny/permit (MAC UDK ACL rule)” on page 922
- the command “deny/permit (IPv4 UDK ACL rule)” on page 924
- the command “deny/permit (IPv4 TCP UDK ACL rule)” on page 926
- the command “deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)” on page 929
- the command “deny/permit (IPv4 ICMP UDK ACL rule)” on page 932
- the command “show ipv4 access-lists” on page 940
- the command “show ipv4-udk access-lists” on page 941
- the command “show ipv6 access-lists” on page 942
- the command “show mac access-lists” on page 943
- the command “show mac-udk access-lists” on page 945
- the command “show access-lists policers (ipv4/ipv4-udk/ipv6/mac/mac-udk)” on page 949
- the command “show qos interface mlag-port-channel” on page 1074
- the command “clear ip igmp snooping counters” on page 1036
- the command “dcb priority-flow-control mode” on page 1051
- the command “show buffers status” on page 1132
- the command “storm-control” on page 1142
- the command “switchmode store-and-forward” on page 1208
- Section 6.1.2.3, “Virtual Routing and Forwarding,” on page 1219
- the command “vrf definition” on page 1223
- the command “show vrf” on page 1233
- the command “show ip interface vrf” on page 1259
- the command “ip route” on page 1274
- the command “show ip route” on page 1276
- the command “show {ip | ipv6} bgp” on page 1435
- the command “show ip pim rp-candidate” on page 1531
- the command “ip mroute” on page 1535

- the command “ip dhcp relay instance (interface config)” on page 1590
- the command “show ip dhcp relay” on page 1593
- the command “ipv6 dhcp relay instance” on page 1596
- the command “show ipv6 dhcp relay” on page 1605
- Appendix B, “Show Commands Not Supported by JSON” on page 1943

Rev 4.80

Software Ver. 3.6.5000 – November 05, 2017

Added:

- Section 2.1.1, “Configuring the Switch with ZTP,” on page 49
- Section 2.4, “Zero-touch Provisioning,” on page 53
- the command “logging level” on page 349
- the command “show log” on page 355
- Section 4.10.2, “Configuring Signal Degradation Monitoring,” on page 377
- Section 4.21, “Control Plane Policing,” on page 664
- Section 5.11.3, “ACL Logging,” on page 891
- Section 5.11.4, “ACL Capability Summary,” on page 892
- the command “bind-point rif” on page 896
- the command “remark” on page 897
- the command “shared-counter” on page 898
- the command “clear shared-counters” on page 899
- the command “clear counters” on page 900
- the command “{ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list clear counters” on page 901
- the command “deny/permit (IPv6 ACL rule)” on page 914
- the command “deny/permit (IPv6 TCP ACL rule)” on page 916
- the command “deny/permit (IPv6 TCP-UDP/UDP ACL rule)” on page 918
- the command “deny/permit (IPv6 ICMPv6 ACL rule)” on page 920
- the command “deny/permit (MAC UDK ACL rule)” on page 922
- the command “deny/permit (IPv4 UDK ACL rule)” on page 924
- the command “deny/permit (IPv4 TCP UDK ACL rule)” on page 926
- the command “deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)” on page 929
- the command “deny/permit (IPv4 ICMP UDK ACL rule)” on page 932
- the command “port access-group (IPv4/IPv4 UDK/IPv6/MAC/MAC UDK)” on page 934
- the command “access-list log” on page 936

- the command “show access-lists shared-counters (ipv4/ipv4-udk/ipv6/mac/mac-udk)” on page 950
- the command “show access-lists log” on page 952
- the command “show access-lists log config” on page 953
- the command “traffic pool” on page 1114
- the command “type” on page 1115
- the command “map switch-priority” on page 1116
- the command “type map switch-priority” on page 1117
- the command “memory percent” on page 1118
- the command “advanced buffer management” on page 1119
- the command “reserved shared size” on page 1122
- the command “pool size type” on page 1123
- the command “cable-length” on page 1130
- the command “show buffers mode” on page 1131
- the command “show buffers pools” on page 1137
- the command “show buffers pools mc-buffers” on page 1138
- the command “show qos” on page 1072
- the command “show qos interface l2-mapping” on page 1078
- the command “show qos interface l3-mapping” on page 1079
- the command “show qos interface rewrite-mapping” on page 1080
- the command “show qos interface tc-mapping” on page 1081
- the command “show qos mapping ingress interface egress interface” on page 1082
- the command “show traffic pool” on page 1139
- the command “show traffic pool” on page 1140
- Section 5.19, “Head-of-Queue Lifetime Limit,” on page 1145
- Section 5.20, “User Defined Keys,” on page 1147
- the command “clear ip dhcp relay counters” on page 1591
- the command “clear ipv6 dhcp relay counters” on page 1604
- the command “show magp interface vlan” on page 1580

Deleted:

- the command “mc-unaware tc binding”

Updated:

- Table 20, “WebUI IP Route Submenus,” on page 84
- the command “cli max-sessions” on page 91
- the command “show ip dhcp” on page 182
- the command “show interfaces ethernet” on page 730
- the command “show interfaces counters” on page 726

- the command “show isolation-group” on page 750
- the command “show interfaces port-channel” on page 767
- the command “show spanning-tree detail” on page 835
- the command “show spanning-tree vlan” on page 839
- the command “show interfaces mlag-port-channel” on page 876
- the command “{ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list” on page 895
- the command “deny/permit (IPv4 ACL rule)” on page 905
- the command “show access-lists summary” on page 951
- the command “show dcb ets” on page 1094
- the command “map pool type reserved” on page 1125
- the command “pool mc-buffer” on page 1128
- the command “pool description” on page 1129
- the command “show buffers status” on page 1132
- the command “show buffers details” on page 1134
- the command “destination interface” on page 1161
- the command “show monitor session summary” on page 1168
- the command “show ip interface vrf” on page 1259
- the command “show ip route summary” on page 1284
- the command “show ip route interface” on page 1285
- the command “show vrrp detail” on page 1570
- the command “show vrrp statistics” on page 1571
- the command “show magp” on page 1579
- the command “show ip dhcp relay counters” on page 1594
- the command “show ipv6 dhcp relay” on page 1605
- Appendix B, “Show Commands Not Supported by JSON” on page 1943

About this Manual

This manual provides general information concerning the scope and organization of this User's Manual.

Intended Audience

This manual is intended for network administrators who are responsible for configuring and managing Mellanox Technologies' switch platforms.

Related Documentation

The following table lists the documents referenced in this *User's Manual.*

Table 1 - Reference Documents

Document Name	Description
System Hardware User Manual	This document contains hardware descriptions, LED assignments and hardware specifications among other things.
Switch Product Release Notes	Please look up the relevant switch system/series release note file
Mellanox Virtual Modular Switch Reference Guide	This reference architecture provides general information concerning Mellanox L2 and L3 Virtual Modular Switch (VMS) configuration and design.
Onyx XML API Reference Guide	This manual provides general information concerning Onyx XML API.

All of these documents can be found on the Mellanox website. They are available either through the product pages or through the support page with a login and password.

Glossary

Table 2 - Glossary

AAA	Authentication, Authorization, and Accounting. Authentication - verifies user credentials (username and password). Authorization - grants or refuses privileges to a user/client for accessing specific services. Accounting - tracks network resources consumption by users.
ARP	Address Resolution Protocol. A protocol that translates IP addresses into MAC addresses for communication over a local area network (LAN).
CLI	Command Line Interface. A user interface in which you type commands at the prompt
DCB	Data Center Bridging

Table 2 - Glossary

DCBX	DCBX protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX end points exchange request and acknowledgment messages. For flexibility, parameters are coded in a type-length-value (TLV) format.
DHCP	The Dynamic Host Configuration Protocol (DHCP) is an automatic configuration protocol used on IP networks.
DNS	Domain Name System. A hierarchical naming system for devices in a computer network
ETS	ETS provides a common management framework for assignment of bandwidth to traffic classes.
FTP/TFTP/sFTP	File Transfer Protocol (FTP) is a standard network protocol used to transfer files from one host to another over a TCP-based network, such as the Internet.
Gateway	A network node that interfaces with another network using a different network protocol
HA (High Availability)	A system design protocol that provides redundancy of system components, thus enables overcoming single or multiple failures in minimal downtime
Host	A computer platform executing an Operating System which may control one or more network adapters
LACP	Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).
LDAP	The Lightweight Directory Access Protocol is an application protocol for reading and editing directories over an IP network.
LLDP (Link Layer Discovery Protocol)	A vendor neutral link layer protocol used by network devices to advertise their identify, capabilities and for neighbor discovery
MAC	A Media Access Control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used for numerous network technologies and most IEEE 802 network technologies including Ethernet.
MTU (Maximum Transfer Unit)	The maximum size of a packet payload (not including headers) that can be sent /received from a port
Network Adapter	A hardware device that allows for communication between computers in a network
PFC/FC	Priority Based Flow Control applies pause functionality to traffic classes OR classes of service on the Ethernet link.
RADIUS	Remote Authentication Dial In User Service. A networking protocol that enables AAA centralized management for computers to connect and use a network service.

Table 2 - Glossary

RDMA (Remote Direct Memory Access)	Accessing memory in a remote side without involvement of the remote CPU
RSTP	Rapid Spanning Tree Protocol. A spanning-tree protocol used to prevent loops in bridge configurations. RSTP is not aware of VLANs and blocks ports at the physical level.
SA (Subnet Administrator)	The interface for querying and manipulating subnet management data
SCP	Secure Copy or SCP is a means of securely transferring computer files between a local and a remote host or between two remote hosts. It is based on the Secure Shell (SSH) protocol.
SNMP	Simple Network Management Protocol. A network protocol for the management of a network and the monitoring of network devices and their functions
NTP	Network Time Protocol. A protocol for synchronizing computer clocks in a network
SSH	Secure Shell. A protocol (program) for securely logging in to and running programs on remote machines across a network. The program authenticates access to the remote machine and encrypts the transferred information through the connection.
syslog	A standard for forwarding log messages in an IP network
TACACS+	Terminal Access Controller Access-Control System Plus. A networking protocol that enables access to a network of devices via one or more centralized servers. TACACS+ provides separate AAA services.
XML Gateway	Extensible Markup Language Gateway. Provides an XML request-response protocol for setting and retrieving HW management information.

1 Introduction

Mellanox® Onyx™ enables the management and configuration of Mellanox Technologies' switch system platforms.

Onyx provides a full suite of management options, including support for SNMPv1, 2, 3, and web user interface (WebUI). In addition, it incorporates a familiar industry-standard CLI, which enables administrators to easily configure and manage the system.

1.1 System Features

Table 3 - General System Features

Feature	Detail
Software management	• Dual software image • Software and firmware updates
File management	• FTP • TFTP • SCP
Logging	• Event history log • SysLog support
Management interface	• DHCP/Zeroconf • IPv6
Chassis management	• Monitoring environmental controls • Power management • Auto-temperature control • High availability
Network management interfaces	• SNMP v1,v2c,v3 • interfaces (XML Gateway) • Puppet Agent
Security	• SSH • Telnet • RADIUS • TACACS+
Date and time	• NTP
Cables & transceivers	• Transceiver info

1.2 Ethernet Features

Table 4 - Ethernet Features

Feature	Detail
General	• ACL – 6400 rules (permit/deny) • Breakout cables • Jumbo Frames (9K)
Ethernet support	• 90100 unicast MAC addresses • DCBX • DHCP Relay • ETS (802.1Qaz) • Flow control (802.3x) • IGMP snooping v1,2 • LAG/LACP (802.3ad), 16 links per LAG (64 LAGs) • LLDP • MLAG • MSTP • OpenFlow 1.3 • PFC (802.1Qbb) • Rapid Spanning Tree (802.1w) • sFlow • VLAN (802.1Q) – 4K
IP routing	• 50K ARP entries • BGP • DHCP Relay • ECMP • IGMP • IPv4 • IPv6 • OSPF • PIM • VLAN interface • Loopback interface • Router interface • VRRP

Figure 1: Managing an Ethernet Fabric Using Onyx

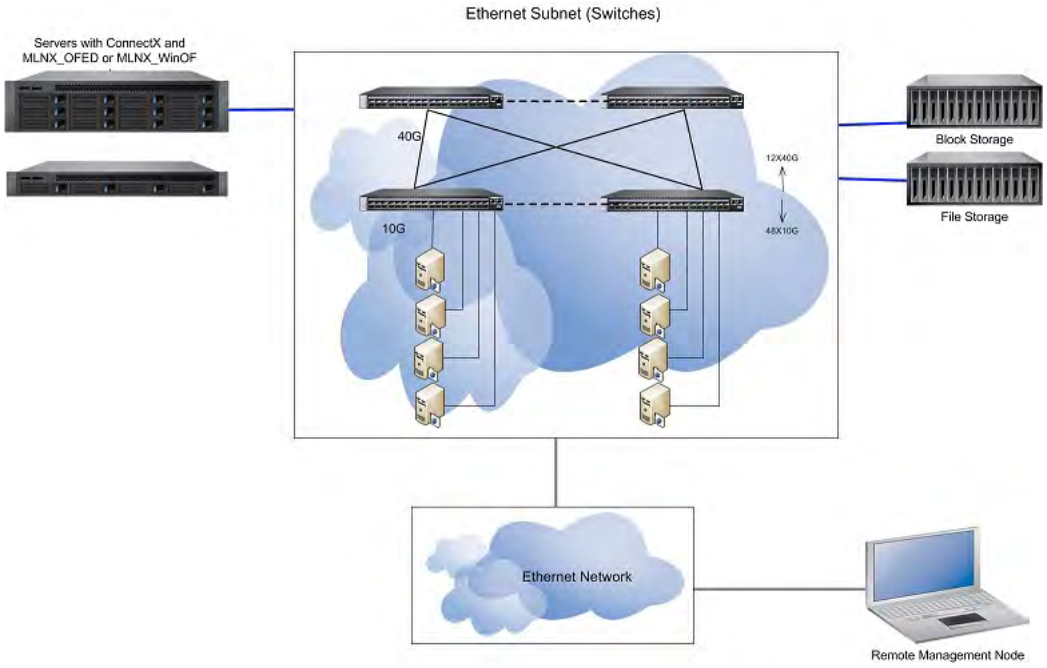

2 Getting Started

The procedures described in this chapter assume that you have already installed and powered on your switch according to the instructions in the *Hardware Installation Guide*, which was shipped with the product.

2.1 Configuring the Switch for the First Time

➢ ***To configure the switch:***

Step 1. Connect the host PC to the console (RJ-45) port of the switch system using the supplied cable. The console ports for systems are shown below.

Make sure to connect to the console RJ-45 port of the switch and not to the MGT port.

DHCP is enabled by default over the MGT port. Therefore, if you have configured your DHCP server and connected an RJ-45 cable to the MGT port, simply log in using the designated IP address.

Step 2. Configure a serial terminal with the settings described below.

This step may be skipped if the DHCP option is used and an IP is already configured for the MGT port.

Table 5 - Serial Terminal Program Configuration

Parameter	Setting
Baud Rate	115200
Data bits	8
Stop bits	1
Parity	None
Flow Control	None

Step 3. You are prompted with the boot menu.

```
Mellanox  Boot Menu:

   1: <image #1>
   2: <image #2>
   u: USB menu (if USB device is connected) (password required)
   c: Command prompt (password required)

   Choice:
```

Select “1” to boot with software version installed on partition #1.
Select “2” to boot with software version installed on partition #2.
Selecting “u” is not currently supported.

The Onyx Boot Menu features a countdown timer. It is recommended to allow the timer to run out by not selecting any of the options.

Step 4. Log in as *admin* and use *admin* as password.

If the machine is still initializing, you might not be able to access the CLI until initialization completes. As an indication that initialization is ongoing, a countdown of the number of remaining modules to be configured is displayed in the following format: “<no. of modules> Modules are being configured”.

Step 5. Go through the configuration wizard.

The following table shows an example of a wizard session.

Table 6 - Configuration Wizard Session - IP Configuration by DHCP (Sheet 1 of 2)

Wizard Session Display (Example)	Comments
Mellanox configuration wizard Do you want to use the wizard for initial configuration? yes	You must perform this configuration the first time you operate the switch or after resetting the switch to the factory defaults. Type “y” and then press <Enter>.
Step1: Hostname? [switch-1]	If you wish to accept the default hostname, then press <Enter>. Otherwise, type a different hostname and press <Enter>.
Step 2: Use DHCP on mgmt0 interface? [yes]	Perform this step to obtain an IP address for the switch. (mgmt0 is the management port of the switch.) If you wish the DHCP server to assign the IP address, type “yes” and press <Enter>. If you type “no” (no DHCP), then you will be asked whether you wish to use the “zeroconf” configuration or not. If you enter “yes” (yes Zeroconf), the session will continue as shown in Table 7. If you enter “no” (no Zeroconf), then you need to enter a *static* IP, and the session will continue as shown in Table 8.
Step 3: Enable IPv6 [yes]	Perform this step to enable IPv6 on management ports. If you wish to enable IPv6, type “yes” and press <Enter>. If you enter “no” (no IPv6), then you will automatically be referred to Step 5.

Table 6 - Configuration Wizard Session - IP Configuration by DHCP (Sheet 2 of 2)

Wizard Session Display (Example)	Comments
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface	Perform this step to enable StateLess address autoconfig on external management port. If you wish to enable it, type “yes” and press <Enter>. If you wish to disable it, enter “no”.
Step 5: Use DHCPv6 on mgmt0 interface? [yes]	Perform this step to enable DHCPv6 on the MGMT0 interface.
Step 5: Admin password (Press <Enter> to leave unchanged)? <new_password> Step 4: Confirm admin password? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. Then confirm the password by re-entering it. Note that password characters are *not* printed.
You have entered the following information: 1. Hostname: <switch name> 2. Use DHCP on mgmt0 interface: yes 3. Enable IPv6: yes 4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes 5. Enable DHCPv6 on mgmt0 interface: no 6. Admin password (Enter to leave unchanged): (CHANGED) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: <Enter> Configuration changes saved. To return to the wizard from the CLI, enter the “configuration jump-start” command from configuration mode. Launching CLI... <switch name> [standalone: master] >	The wizard displays a summary of your choices and then asks you to confirm the choices or to re-edit them. Either press <Enter> to save changes and exit, or enter the configuration step number that you wish to return to. Note: To run the command “configuration jump-start” you must be in Config mode.

Table 7 - Configuration Wizard Session - IP Zeroconf Configuration

Wizard Session Display - IP Zeroconf Configuration (Example)
Mellanox configuration wizard Do you want to use the wizard for initial configuration? y Step 1: Hostname? [switch-112126] Step 2: Use DHCP on mgmt0 interface? [no] Step 3: Use zeroconf on mgmt0 interface? [no] yes Step 4: Default gateway? [192.168.10.1] Step 5: Primary DNS server? Step 6: Domain name? Step 7: Enable IPv6? [yes] yes Step 8: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no] no Step 9: Admin password (Enter to leave unchanged)? You have entered the following information: 1. Hostname: switch-112126 2. Use DHCP on mgmt0 interface: no 3. Use zeroconf on mgmt0 interface: yes 4. Default gateway: 192.168.10.1 5. Primary DNS server: 6. Domain name: 7. Enable IPv6: yes 8. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes 9. Admin password (Enter to leave unchanged): (unchanged) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: Configuration changes saved. To return to the wizard from the CLI, enter the “configuration jump-start” command from configure mode. Launching CLI... <switch name> [standalone: master] >

Table 8 - Configuration Wizard Session - Static IP Configuration

Wizard Session Display - Static IP Configuration (Example)
Mellanox configuration wizard Do you want to use the wizard for initial configuration? y Step 1: Hostname? [switch-112126] Step 2: Use DHCP on mgmt0 interface? [yes] n Step 3: Use zeroconf on mgmt0 interface? [no] Step 4: Primary IP address? 192.168.10.4 Mask length may not be zero if address is not zero (interface mgmt0) Step 5: Netmask? [0.0.0.0] 255.255.255.0 Step 6: Default gateway? 192.168.10.1 Step 7: Primary DNS server? Step 8: Domain name? Step 9: Enable IPv6? [yes] yes Step 10: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no] no Step 11: Admin password (Enter to leave unchanged)? You have entered the following information: 1. Hostname: switch-112126 2. Use DHCP on mgmt0 interface: no 3. Use zeroconf on mgmt0 interface: no 4. Primary IP address: 192.168.10.4 5. Netmask: 255.255.255.0 6. Default gateway: 192.168.10.1 7. Primary DNS server: 8. Domain name: 9. Enable IPv6: yes 10. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: no 11. Admin password (Enter to leave unchanged): (unchanged) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: Configuration changes saved. To return to the wizard from the CLI, enter the “configuration jump-start” command from configure mode. Launching CLI... <switch name>[standalone: master] >

Step 6. Check the mgmt0 interface configuration before attempting a remote (for example, SSH) connection to the switch. Specifically, verify the existence of an IP address.

```
switch # show interfaces mgmt0

Interface mgmt0 status:
  Comment         :
  Admin up        : yes
  Link up         : yes
  DHCP running    : yes
  IP address      : 10.12.67.34
  Netmask         : 255.255.0.0
  IPv6 enabled    : yes
  Autoconf enabled: no
  Autoconf route  : yes
  Autoconf privacy: no
  DHCPv6 running  : no
  IPv6 addresses  : 1

  IPv6 address:
    fe80::268a:7ff:fe53:3d8e/64

  Speed           : 1000Mb/s (auto)
  Duplex          : full (auto)
  Interface type  : ethernet
  Interface source: physical
  MTU             : 1500
  HW address      : 00:02:C9:11:A1:B2

  Rx:
   11700449  bytes
      55753  packets
          0  mcast packets
          0  discards
          0  errors
          0  overruns
          0  frame

  Tx:
   5139846  bytes
     28452  packets
         0  discards
         0  errors
         0  overruns
         0  carrier
         0  collisions
      1000  queue len
```

2.1.1 Configuring the Switch with ZTP

Mellanox Onyx™ Zero-touch Provisioning (ZTP) automates initial configuration of switch systems at boot time. It helps minimize manual operation and reduce customer initial deployment cost.

For more information, please refer to Section 2.4, “Zero-touch Provisioning,” on page 53.

2.1.2 Rerunning the Wizard

➢ ***To rerun the wizard:***

Step 1. Enter the config mode:

```
switch > enable
switch # config terminal
```

Step 2. Rerun the wizard:

```
switch (config) # configuration jump-start
```

2.2 Starting the Command Line (CLI)

Step 1. Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.

Step 2. Start a remote secured shell (SSH) to the switch using the command “ssh -l <username> <switch ip address>.”

```
rem_mach1 >  ssh -l <username> <ip address>
```

Step 3. Log into the switch (default username is *admin*, password *admin*)

Step 4. Read and accept the EULA when prompted.

Step 5. Once you get the prompt, you are ready to use the system.

```
Mellanox Onyx Switch Management

Password:
Last login: <time> from <ip-address>

Mellanox Switch
Please read and accept the Mellanox End User License Agreement located at:
https://www.mellanox.com/related-docs/prod_management_software/MLNX_Onyx_EULA.pdf

switch >
```

2.3 Starting the Web User Interface (WebUI)

➢ ***To start a WebUI connection to the switch platform:***

WebUI access is enabled by default.
To disable web access, run the command “no web http enable” or “no web https enable” through the CLI.

Step 1. Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.

Step 2. Open a web browser – Firefox 12, Chrome 18, IE 8, Safari 5 or higher.

Note: Make sure the screen resolution is set to 1024*768 or higher.

Step 3. Type in the IP address of the switch or its DNS name in the format: https://<switch_IP_address>.

Step 4. Log into the switch (default user name is *admin*, password *admin*).

Figure 2: Onyx Login Window

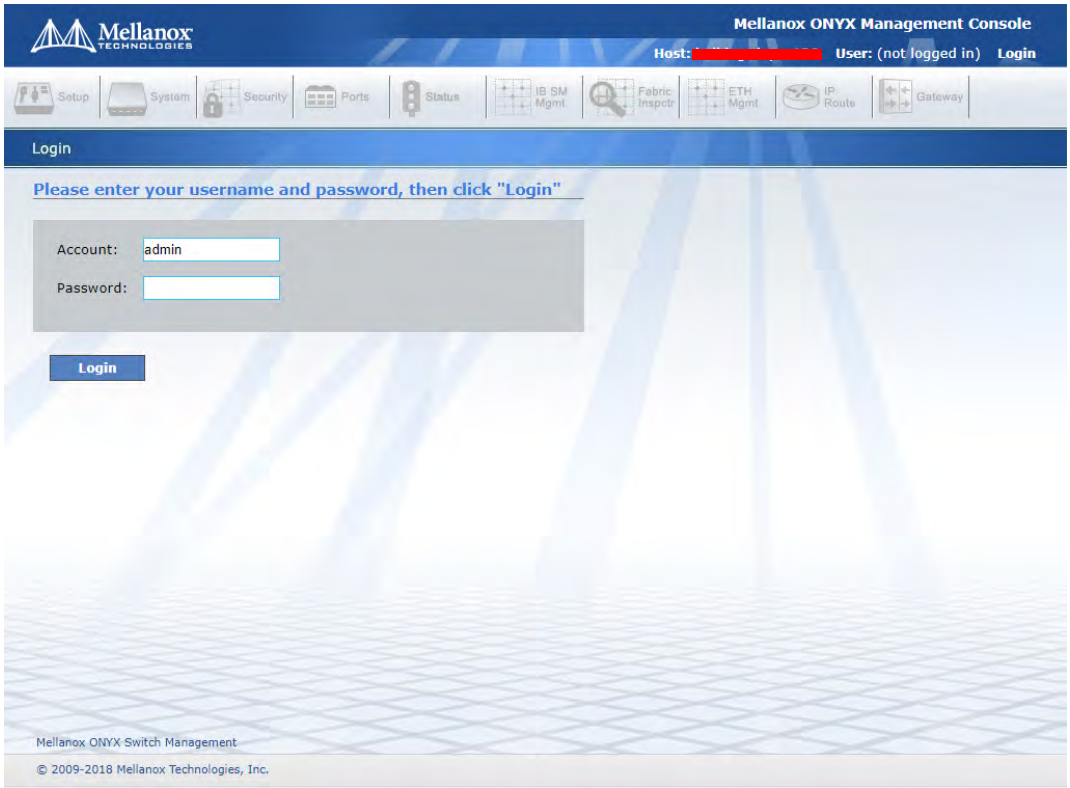

Step 5. Read and accept the EULA if prompted.
You are only prompted if you have not accessed the switch via CLI before.

Step 6. The Welcome popup appears. After reading through the content, click OK to continue.
You may click on the links under Documentation to reach the Onyx documentation.

The link under What's New takes you straight to the Changes and New Features section of the switch OS Release Notes.

Figure 3: Welcome Popup (Example)

You may also tick the box to not show this popup again. But should you wish to see this window again, click "Product Documents" on the upper right corner of the WebUI.

Step 7. A default status summary is displayed as shown in Figure 4.

Figure 4: Display After Login

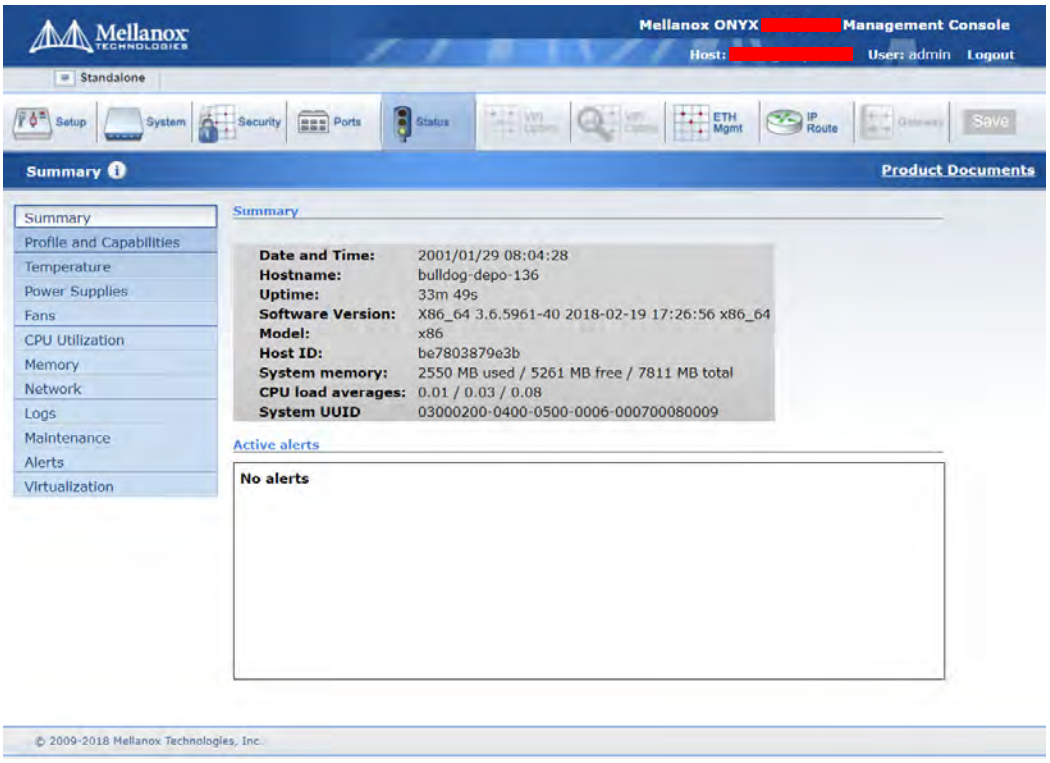

2.4 Zero-touch Provisioning

Onyx Zero-touch Provisioning (ZTP) automates initial configuration of Mellanox switches at boot time. It helps minimize manual operation and reduce customer initial deployment cost. Onyx ZTP allows the customer to automatically upgrade the switch with a specified OS image, set up initial configuration database, and load and run a container image file.

The initial configuration is applied using a regular text file. The user can create such a configuration file by editing the output of a "`show running-config`" command.

Only a textual configuration files is supported.

The user-defined docker image can be used by customers to run their own applications in a sandbox on an Onyx platform. And can therefore be also used for automating initial configuration.

Only one docker container could be launched in ZTP.

2.4.1 Running DHCP-ZTP

There is no explicit command to enable ZTP. It is enabled by default. Disabling it is performed by a user-initiated configuration save (using the command "`configuration write`"). The only way to re-enable ZTP would be to run a "`reset factory`" command, clearing the configuration of the switch and rebooting the system.

Onyx ZTP is based on DHCP. For ZTP to work, Onyx enables DHCP by default on all its management interfaces. Onyx requests option 66 (tftp-server-name) and 67 (bootfile-name) from the DHCPv4 server or option 58 (bootfile-url) from the DHCPv6 server, and waits for the DHCP responses which contain file URLs. The DHCP server must be configured to send back the URLs for the software image, configuration file, and docker container image via these two options. Option 66 would contain the URL prefix to the location of the files, option 67 would contain the name of files, and option 58 would contain the complete URLs of files. The format of these two options is a string list separated by commas. The list items are placed in a fixed order:

```
<image file>, <config file>, <docker container file>
```

The item value can be empty, but the comma shall not be omitted.

To have DHCP server figure out the proper files based on switch specific information, Onyx must provide some sort of identity information for the server to classify the switches. Besides the aforementioned options, Onyx attaches option 43 (vendor specific information) and option 60 (vendor class identifier) in DHCPv4 requests, and option 17 (vendor-opts) in DHCPv6. Option 60 is set as string "Mellanox" and options 17 and 43 contain the following Mellanox-specific sub-options:

- System Model
- Chassis Part Number

- Chassis Serial Number
- Management MAC
- System Profile
- Onyx Release Version

The corresponding subtypes respectively are defined as:

```
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_MODEL        1
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_PARTNUM      2
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_SERIAL       3
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_MAC          4
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_PROFILE      5
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_RELEASE      6
```

Upon receiving such DHCP requests from a client, the server should be able to map the switch-specific information to the target file URLs according to predefined rules.

Once Onyx receives the URLs from the DHCP server, it executes ZTP as follows:

1. If the software image URL is not specified, this step is skipped. Otherwise:
   - Perform disk space cleanup if necessary and fetch the image if it does not exist locally
   - Resolve the image version:
     - If it is already installed on active partition, proceed to step 2
     - If it is installed on a standby partition, switch partition and reboot
     - If it is not installed locally, install it and switch to the new image and then reboot
     - In case of reboot, ZTP performs step 1 again and no image upgrade will occur
2. If configuration file URL is not specified, skip this step. Otherwise:
   - Fetch the configuration file
   - Apply the configuration file
3. Skip these steps if a docker image file URL is not specified. Otherwise:
   - Fetch the docker image file
   - Load the docker image
   - Clean up the docker images with the same name and different tag.
   - Start the container based on the image
   - Remove the downloaded docker image file

While performing file transfer via HTTP, the same information as DHCP option 43 is expected to be carried in a HTTP GET request. Onyx supports the following proprietary HTTP headers:

- MlnxSysProfile
- MlnxMgmtMac
- MlnxSerialNumber
- MlnxModelName
- MlnxPartNumber
- MlnxReleaseVersion

In case of failure, the switch waits a random number of seconds between 1 and 20 and reattempts the operation. The switch attempts this up to 10 times.

ZTP progress is printed to terminals including console and active SSH sessions.

2.4.2 ZTP on Director Switches

For director switch systems, the two management nodes start ZTP individually. Status synchronization is then performed between the two nodes:

- Target software image version needs to be the same, otherwise ZTP fails
- Both nodes must install the software image successfully, otherwise ZTP fails
- ZTP failure for one node leads to failure for both
- ZTP disable on one node leads to ZTP disable for both
- ZTP abort on one node leads to ZTP abort for both

In ZTP configuration files, commands between #<CHASSIS_MASTER> and #</CHASSIS_MASTER> pair are only executed on the master.

```
#<CHASSIS_MASTER>
   chassis ha bip 10.7.146.34 /24
#</CHASSIS_MASTER>
```

Node reboot caused by ZTP is also synchronized:

1. Master node asks slave to reboot.
2. Slave node switches to next boot location and acknowledges the reboot request.
3. Master node reboots slave node via hardware.
4. Master node reboots itself.

2.4.3 ZTP and Onyx Software Upgrade

Software upgrade from non-ZTP versions to ZTP versions and vice versa is supported. When upgrading from a non-ZTP version, ZTP is disabled because ZTP is always assumed to start with an empty configuration, otherwise the final configuration becomes a mixture of the existing configuration from the stored database and new configuration from the server and hence not deterministic.

2.4.4 DHCPv4 Configuration Example

The following is a URL configuration example for ISC DHCPv4 server:

```
 host master {
          hardware ethernet E4:1D:2D:5B:72:80;
          fixed-address 3.1.2.13;
          option tftp-server-name "scp://<user>:<password>@3.1.3.100/ztp/,scp://
                                   <user>:<password>@3.1.3.100/ztp/,scp://
                                   <user>:<password>@3.1.3.100/ztp/";
          option bootfile-name "image-X86_64-3.6.4612.img, switch-1.conf,
ubuntu.img.gz";
     }
```

DHCPv4 request is made out of the following components:

- Option 43 (vendor-encapsulated-options) and option 60 (vendor-class-identifier) are added in the DHCPv4 request packet
- Option 66 (tftp-server-name) and option 67 (bootfile-name) are added in the parameter request list of DHCPv4 request packet

2.4.5 DHCPv6 Configuration Example

The following is a DHCPv6 configuration example:

```
 host master {
         ......
        option dhcp6.bootfile-url "scp://<user>:<password>@[2000::1]/ztp/image-X86_64-
                                   3.6.4612.img, scp://<user>:<password>@[2000::1]/ztp/
                                   switch.conf, scp://<user>:<password>@[2000::1]/ztp/
                                   ubuntu.img.gz";
     }
```

DHCPv6 request is made out of the following components:

- Option 17 (vendor-opts) is added in the DHCPv6 request packet
- Option 59 (bootfile-url) is added in the parameter request list of DHCPv6 request packet

2.4.6 Commands

no zero-touch suppress-write

no zero-touch suppress-write

The no form of the command disables suppression of configuration write.

Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.5000
Role	admin
Example	`switch (config) # no zero-touch suppress-write`
Related Commands	show zero-touch
Notes	When ZTP is active, “configuration write” is suppressed because it may interfere with ZTP operation. Therefore, after running “no zero-touch suppress-write” if “configuration write” is performed, then ZTP is disabled as a consequence of the database save.

zero-touch abort

	zero-touch abort Aborts on-going zero-touch process.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.5000
Role	admin

Example	`switch (config) # zero-touch abort`
	
	`Zero-touch failed [Zero-touch is aborted by operator]`
	`Zero-touch provisioning will be aborted`

Related Commands	show zero-touch
Notes	

show zero-touch

	show zero-touch Displays zero-touch status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin

Example	`switch (config) # show zero-touch`
	`Zero-Touch status:`
	`  Active: yes`
	`  Status: Waiting for zero-touch start`
	`  Suppress-write: no`
	`  Configured by zero-touch: no`
	`  Configuration changed after zero-touch: no`

Related Commands	zero-touch abort zero-touch suppress-write
Notes	

2.5 Licenses

Onyx software package can be extended with premium features. Installing a license allows you to access the specified premium features.

This section is relevant only to switch systems with an internal management capability.

2.5.1 Installing Onyx License (CLI)

➢ ***To install an Onyx license via CLI:***

Step 1. Before applying a license, please make sure your system's time is configured correctly by manually setting it using the CLI command "clock set", or by using NTP using the command "ntp".

Step 2. Login as *admin* and change to *Config* mode.

```
switch > enable
switch # config terminal
```

Step 3. Install the license using the key. Run:

```
switch (config) # license install <license key>
```

Step 4. Display the installed license(s) using the following command.

```
switch (config) # show licenses
License 1: <license key>
Feature: EFM_SX
Valid: yes
Active: yes
switch (config) #
```

Make sure that the "Valid" and "Active" fields both indicate "yes".

Step 5. Save the configuration to complete the license installation. Run:

```
switch (config) # configuration write
```

If you do not save the installation session, you will lose the license at the next system start up.

2.5.2 Installing Onyx License (Web)

➢ ***To install an Onyx license via WebUI:***

Step 1. Log in as *admin*.

Step 2. Click the **Setup** tab and then **Licensing** on the left side navigation pane.

Figure 5: No Licenses Installed

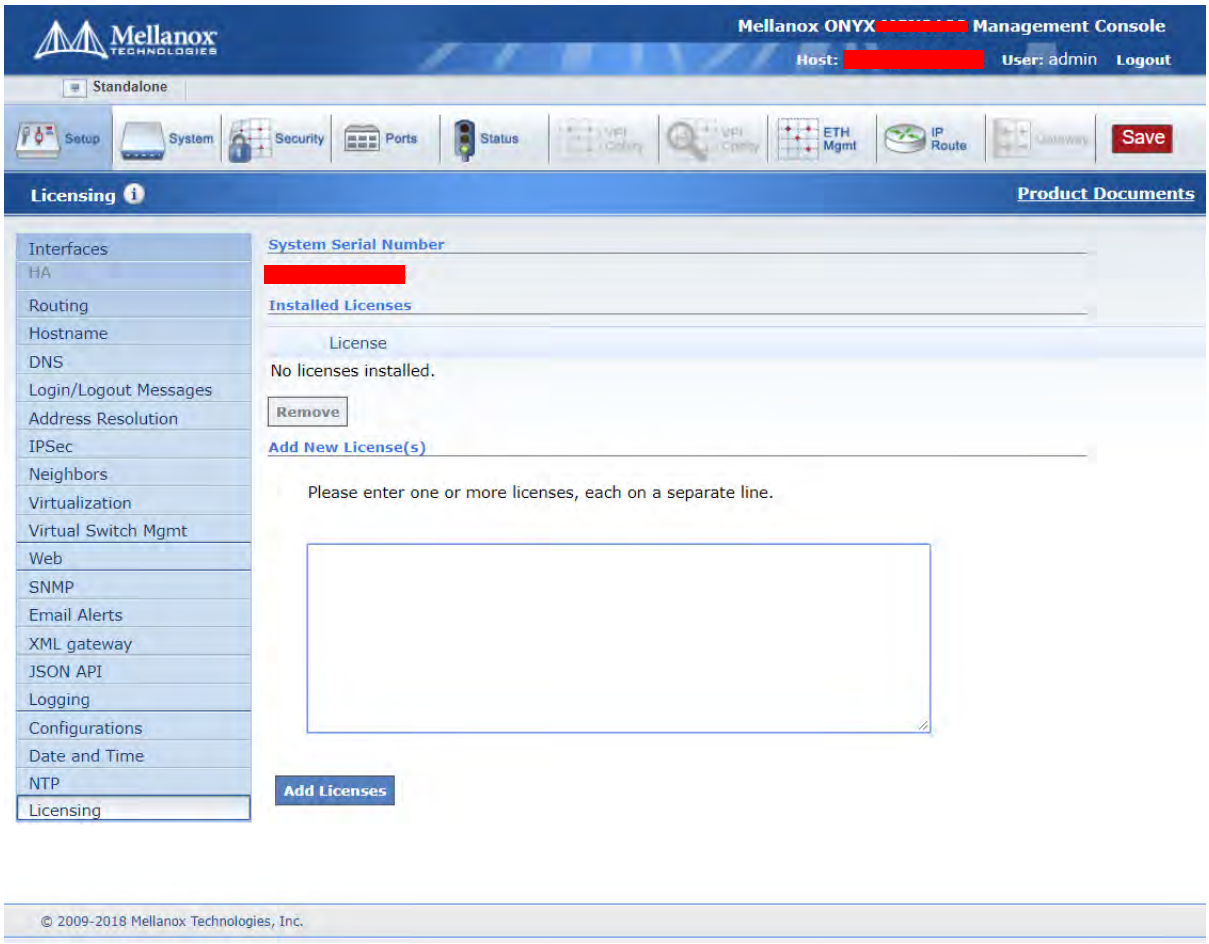

Step 3. Enter your license key(s) in the text box. If you have more than one license, please enter each license in a separate line. Click "Add Licenses" after entering the last license key to install them.

If you wish to add another license key in the future, you can simply enter it in the text box and click "Add Licenses" to install it.

Figure 6: Enter License Key(s) in Text Box

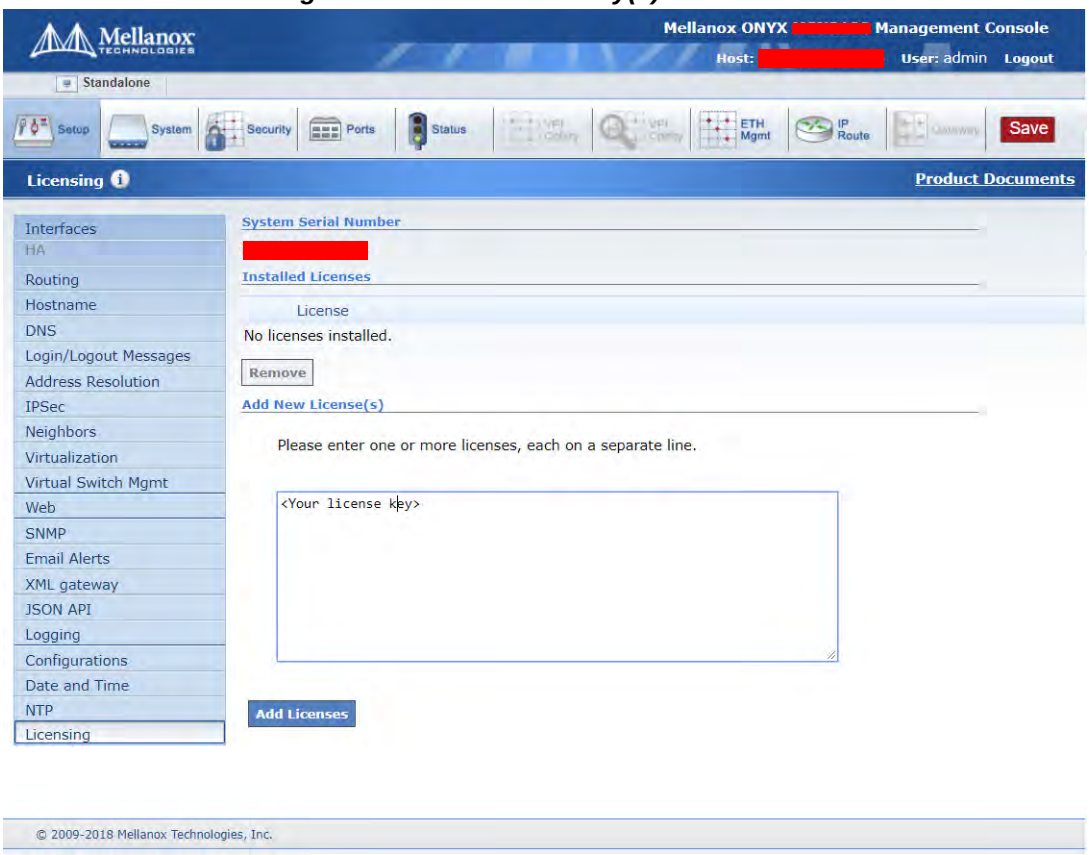

All installed licenses should now be displayed.

Figure 7: Installed License

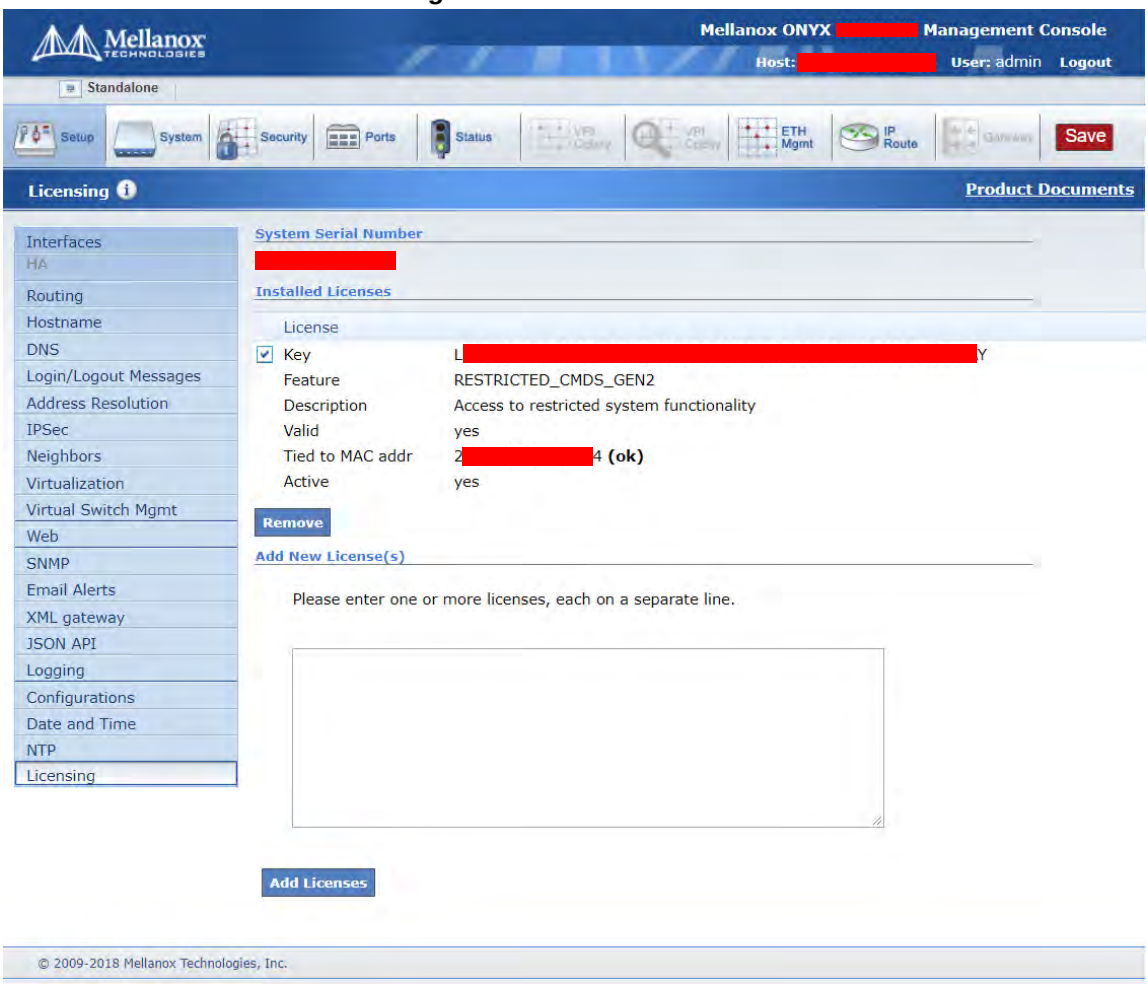

Step 4. Save the configuration to complete the license installation.

If you do not save the installation session, you will lose the installed licenses at the next system boot.

2.5.3 Retrieving a Lost License Key

In case of a lost Onyx license key, contact your authorized Mellanox reseller and provide the switch's chassis serial number.

➢ ***To obtain the switch's chassis serial number:***

Step 1. Login to the switch.

Step 2. Retrieve the switch's *chassis serial number* using the command "show inventory".

```
switch (config) # show inventory
---------------------------------------------------------------------------
Module           Part Number        Serial Number        Asic Rev.    HW Rev.
---------------------------------------------------------------------------
CHASSIS          MSN2100-CB2F       MT1752X06330         N/A          B3
MGMT             MSN2100-CB2F       MT1752X06330         1            B3
```

Step 3. Send your Mellanox reseller the following information to obtain the license key:

- The chassis serial number
- The type of license you need to retrieve. Refer to "Licenses" on page 60.

Step 4. Once you receive the license key, you can install the license as described in the sections above.

2.5.4 Commands

file eula upload

file eula upload <filename> <URL>

Uploads the Mellanox End User License Agreement to a specified remote location.

Syntax Description	filename	The Mellanox End User License Agreement
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
Role	monitor/admin	
Example	`switch (config) # file help-docs upload Mellanox_End_User_License_Agreement.pdf <scp://username[:password]@hostname/path/filename>` `switch (config) #`	
Related Commands	license	
Note		

file help-docs upload

file help-docs upload <filename> <URL or scp://username[:password]@host-name/path/filename>

Uploads the Onyx UM or RN to a specified remote location.

Syntax Description	filename	The file to upload to a remote host
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
Role	admin	
Example	`switch (config) # file help-docs upload Onyx_ETH_User_Manual.pdf <scp://username[:password]@hostname/path/filename>`	
Related Commands		
Note		

license delete

license delete <license-number>

Removes license keys by ID.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Role	admin
Example	`switch (config) # license delete <license-number>`
Related Commands	
Note	Before deleting a license from a switch which is configured to a system profile other than its default, the user must first disable all interfaces and then return the switch to its default system profile.

license install

license install <license-key>

Installs a new license key.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Role	admin
Example	`switch (config) # licenses install <license-key>` `switch (config) #`
Related Commands	
Note	

show licenses

	show licenses Displays a list of all installed licenses. For each license, the following is displayed: • a unique ID which is a small integer • the text of the license key as it was added • whether or not it is valid and active • which feature(s) it is activating • a list of all licensable features specifying whether or not it is currently activated by a license
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Role	admin

Example	`switch (config) # show licenses`
	`License 1: <license key>`
	`Feature: SX_CONFIG`
	`Valid: yes`
	`Active: yes`
	`switch (config) #`

Related Commands	
Note	

3 User Interfaces

3.1 LED Indicators

Table 9 - LED Behavior Details

LED	Qty.	Color	Description
QSFP LEDs	8	Green/Amber	Off – link is down Solid green – link is up Blinking green – data activity. Blinking frequency is proportional to data transfer speed. Blinking amber – link error
Health LED	1	Red/Green/Amber	Off – no power Blinking amber – fault Solid green – normal Solid red – CANMIC boot failure
UID LED	1	Blue	Solid – LED is activated to identify this module

3.2 Command Line Interface Overview

Mellanox Onyx™ is equipped with an industry-standard command line interface (CLI). The CLI is accessed through SSH or Telnet sessions, or directly via the console port on the front panel (if it exists).

3.2.1 CLI Modes

The CLI can be in one of following modes, and each mode makes available a certain group (or level) of commands for execution. The following are some of the CLI configuration modes:

Table 10 - CLI Modes and Config Context

Configuration Mode	Description
Standard	When the CLI is launched, it begins in Standard mode. This is the most restrictive mode and only has commands to query a restricted set of state information. Users cannot take any actions that directly affect the system, nor can they change any configuration.
Enable	The `enable` command moves the user to Enable mode. This mode offers commands to view all state information and take actions like rebooting the system, but it does not allow any configurations to be changed. Its commands are a superset of those in Standard mode.

Table 10 - CLI Modes and Config Context

Configuration Mode	Description
config	The `configure terminal` command moves the user from Enable mode to Config mode. Config mode is allowed only for user accounts in the "admin" role (or capabilities). This mode has a full unrestricted set of commands to view anything, take any action, and change any configuration. Its commands are a super-set of those in Enable mode. To return to Enable mode, enter `exit` or `no configure`. Note that moving directly from/to Standard mode to/from Config mode is not possible.
config interface management	Configuration mode for management interface mgmt0, mgmt1 and loopback
config interface ethernet	Configuration mode for Ethernet interface
config interface port-channel	Configuration mode for Port channel (LAG)
config vlan	Configuration mode for VLAN
Any command mode	Several commands such as "show" can be applied within any context

3.2.2 Syntax Conventions

To help you identify the parts of a CLI command, this section explains conventions of presenting the syntax of commands.

Table 11 - Syntax Conventions

Syntax Convention	Description	Example
< > Angled brackets	Indicate a value/variable that must be replaced.	<1...65535> or <switch inter-face>
[ ] Square brackets	Enclose optional parameters. However, only one parameter out of the list of parameters listed can be used. The user cannot have a combination of the parameters unless stated otherwise.	[destination-ip | destination-port | destination-mac]
{ } Braces	Enclose alternatives or variables that are required for the parameter in square brackets.	[mode {active | on | passive}]
| Vertical bars	Identify mutually exclusive choices.	active | on | passive

Do not type the angled or square brackets, vertical bar, or braces in command lines. This guide uses these symbols only to show the types of entries.

CLI commands and options are in lowercase and are case-sensitive.

For example, when you enter the `enable` command, enter it all in lowercase. It cannot be ENABLE or Enable. Text entries you create are also case-sensitive.

3.2.3 Getting Help

You may request context-sensitive help at any time by pressing "?" on the command line. This will show a list of choices for the word you are on, or a list of top-level commands if you have not typed anything yet.

For example, if you are in Standard mode and you type "?" at the command line, then you will get the following list of available commands.

```
switch > ?
cli              Configure CLI shell options
enable           Enter enable mode
exit             Log out of the CLI
help             View description of the interactive help system
no               Negate or clear certain configuration options
show             Display system configuration or statistics
slogin           Log into another system securely using ssh
switch           Configure switch on system
telnet           Log into another system using telnet
terminal         Set terminal parameters
traceroute       Trace the route packets take to a destination
switch-11a596 [standalone: master] >
```

If you type a legal string and then press "?" *without* a space character before it, then you will either get a description of the command that you have typed so far or the possible command/parameter completions. If you press "?" *after* a space character and "<cr>" is shown, this means that what you have entered so far is a complete command, and that you may press Enter (carriage return) to execute it.

Try the following to get started:

```
?
show ?
show c?
show clock?
show clock ?
show interfaces ?     (from enable mode)
```

You can also enter "help" to view a description of the interactive help system.

Note also that the CLI supports command and/or parameter tab-completions and their shortened forms. For example, you can enter "en" instead of the "enable" command, or "cli cl" instead of "cli clear-history". In case of ambiguity (more than one completion option is available, that is), then you can hit double tabs to obtain the disambiguation options. Thus, if you are in Enable

mode and wish to learn which commands start with the letter “c”, type “c” and click twice on the tab key to get the following:

```
switch # c<tab>
clear       cli       configure
switch # c
```

(There are three commands that start with the letter “c”: `clear`, `cli` and `configure`.)

3.2.4 Prompt and Response Conventions

The prompt always begins with the hostname of the system. What follows depends on what command mode the user is in. To demonstrate by example, assuming the machine name is “`switch`”, the prompts for each of the modes are:

```
switch >              (Standard mode)
switch #              (Enable mode)
switch (config) #     (Config mode)
```

The following session shows how to move between command modes: \

```
switch >                                    (You start in Standard mode)
switch > enable                             (Move to Enable mode)
switch #                                    (You are in Enable mode)
switch # configure terminal                       (Move to Config mode)
switch (config) #                           (You are in Config mode)
switch (config) # exit                      (Exit Config mode)
switch #                                    (You are back in Enable mode)
switch # disable                            (Exit Enable mode)
switch >                                    (You are back in Standard mode)
```

Commands entered do not print any response and simply show the command prompt after you press <Enter>.

If an error is encountered in executing a command, the response will begin with “%”, followed by some text describing the error.

3.2.5 Using the “no” Form

Several Config mode commands offer the negation form using the keyword “no”. This no form can be used to disable a function, to cancel certain command parameters or options, or to reset a parameter value to its default. To re-enable a function or to set cancelled command parameters or options, enter the command without the “no” keyword (with parameter values if necessary).

The following example performs the following:

1. Displays the current CLI session options.
2. Disables auto-logout.
3. Displays the new CLI session options (auto-logout is disabled).
4. Re-enables auto-logout (after 15 minutes).
5. Displays the final CLI session options (auto-logout is enabled)

```
// 1. Display the current CLI session options
switch (config) # show cli
```

```
CLI current session settings:
  Maximum line size:        8192
  Terminal width:           157 columns
  Terminal length:          60 rows
  Terminal type:            xterm
  Auto-logout:              15 minutes
  Paging:                   enabled
  Progress tracking:        enabled
  Prefix modes:             enabled
  ...
// 2. Disable auto-logout
switch (config) # no cli session auto-logout
// 3. Display the new CLI session options
switch-1 [standalone: master] (config) # show cli
CLI current session settings:
  Maximum line size:        8192
  Terminal width:           157 columns
  Terminal length:          60 rows
  Terminal type:            xterm
  Auto-logout:              disabled
  Paging:                   enabled
  Progress tracking:        enabled
  Prefix modes:             enabled
  ...
// 4. Re-enable auto-logout after 15 minutes
switch (config) # cli session auto-logout 15
// 5. Display the final CLI session options
switch (config) # show cli
CLI current session settings:
  Maximum line size:        8192
  Terminal width:           157 columns
  Terminal length:          60 rows
  Terminal type:            xterm
  Auto-logout:              15 minutes
  Paging:                   enabled
  Progress tracking:        enabled
  Prefix modes:             enabled
  ...
```

3.2.6 Parameter Key

This section provides a key to the meaning and format of all of the angle-bracketed parameters in all the commands that are listed in this document.

Table 12 - Angled Brackets Parameter Description

Parameter	Description
<domain>	A domain name, e.g. “mellanox.com”.
<hostname>	A hostname, e.g. “`switch-1`”.
<ifname>	An interface name, e.g. “mgmt0”, “mgmt1”, “lo” (loopback), etc.
<index>	A number to be associated with aliased (secondary) IP addresses.
<IP address>	An IPv4 address, e.g. “192.168.0.1”.
<log level>	A syslog logging severity level. Possible values, from least to most severe, are: “debug”, “info”, “notice”, “warning”, “error”, “crit”, “alert”, “emerg”.
<GUID>	Globally Unique Identifier. A number that uniquely identifies a device or component.
<MAC address>	A MAC address. The segments may be 8 bits or 16 bits at a time, and may be delimited by “:” or “.”. So you could say “11:22:33:44:55:66”, “1122:3344:5566”, “11.22.33.44.55.66”, or “1122.3344.5566”.
<netmask>	A netmask (e.g. “255.255.255.0”) or mask length prefixed with a slash (e.g. “/24”). These two express the same information in different formats.
<network prefix>	An IPv4 network prefix specifying a network. Used in conjunction with a netmask to determine which bits are significant. e.g. “192.168.0.0”.
<regular expression>	An extended regular expression as defined by the “grep” in the man page. (The value you provide here is passed on to “grep -E”.)
<node id>	ID of a node belonging to a cluster. This is a numerical value greater than zero.
<cluster id>	A string specifying the name of a cluster.
<port>	TCP/UDP port number.
<TCP port>	A TCP port number in the full allowable range [0...65535].
<URL>	A normal URL, using any protocol that wget supports, including http, https, ftp, sftp, and tftp; or a pseudo-URL specifying an scp file transfer. The scp pseudo-URL format is scp://username:password@hostname/path/filename. Note that the path is an absolute path. Paths relative to the user's home directory are not currently supported. The implementation of ftp does not support authentication, so use scp or sftp for that. Note also that if you omit the “:password” part, you may be prompted for the password in a follow up prompt, where you can type it securely (without the characters being echoed). This prompt will occur if the “cli default prompt empty-password” setting is true; otherwise, the CLI will assume you do not want any password. If you include the “:” character, this will be taken as an explicit declaration that the password is empty, and you will not be prompted in any case.

3.2.7 CLI Pipeline Operator Commands

3.2.7.1 “include” and “exclude” CLI Filtration Options

The Onyx CLI supports filtering “show” commands to display lines containing or excluding certain phrases or characters. To filter the outputs of the “show” commands use the following format:

```
switch (config) # <show command> | {include | exclude} <extended regular expression>
[<ignore-case>] [next <lines>] [prev <lines>]
```

The filtering parameters are separated from the show command they filter by a pipe character (i.e. “|”). Quotation marks may be used to include or exclude a string including space, and multiple filters can be used simultaneously. For example:

```
switch (config) # <show command> | {include <extended regular expression>} [<ignore-
case>] [next <lines>] [prev <lines>] | exclude <extended regular expression> [<ignore-
case>] [next <lines>] [prev <lines>]]
```

Examples:

```
switch (config) # show asic-version | include SX
MGMT              SX                 9.3.3150

arc-switch14 [standalone: master] (config) # show module | exclude PS
======================
Module    Status
======================
MGMT      ready
FAN1      ready
FAN2      ready

switch (config) # show interfaces | include "Eth|discard pac"
Eth1/1
0 discard packets
0 discard packets
Eth1/2
0 discard packets
0 discard packets
Eth1/3
0 discard packets
0 discard packets
Eth1/4
0 discard packets
0 discard packets
switch (config) # show interfaces | include "Tx" next 5 | exclude broad
Tx
0 packets
0 unicast packets
0 multicast packets
0 bytes
--
```

```
Tx
0 packets
0 unicast packets
0 multicast packets
0 bytes
```

3.2.7.2 “watch” CLI Monitoring Option

Onyx also allows viewing a live feed of the progress of any “show” command by using the “watch” option as follows:

```
switch (config) # <show command> | watch [diff] [interval <1-100 secs>]
```

Running the command as such displays an output of the show command that gets updated at a time interval which may be specified using the “interval” parameter (2 seconds by default).

The “diff” parameter highlights the differences between each iteration of the command. For example running the command “`show power | watch diff interval 1`” yields something similar to the following:

```
-------------------------------------------------------------------------------------
Module  Device           Sensor  Power   Voltage  Current  Capacity  Feed  Status
                                 [Watts] [Volts]  [Amp]    [Watts]
-------------------------------------------------------------------------------------
PS1     power-mon        input   85.00   230.00   0.38     460.00    AC    OK
PS2     power-mon        -       -       -        -        -         -     FAIL

Total power used : 85.00 Watts
Total power capacity : 460.00 Watts
Total power available : 375.00 Watts
Maximum consumed power of all turned on modules: 462.00 Watts
```

With the highlighted black blocks indicating the change that has occurred between one iteration of the command from one second to the next.

To exit “watch” mode, press Ctrl+C.

The “watch” option may also be used in conjunction with the “include” and “exclude” options as follows:

```
switch (config) # <show command> | {include | exclude} <extended regular expression> |
watch [diff] [interval <1-100 secs>]
```

For example:

```
switch (config) # show power | include PS | watch diff interval 1
```

3.2.7.3 “json-print” CLI Option

The OnyxMLNX-OS CLI supports printing “show” commands in JSON syntax.

To print the output of the “show” commands as JSON, use the following format:

```
switch (config) # <show command> | json-print
```

Running the command displays an output of the “show” command in JSON syntax structure instead of its regular format. For example:

```
switch (config) # show system profile
Profile: eth-single-switch
Switch (config) # show system profile | json-print
{
"Profile": "eth-single-switch"
}
```

The “json-print” option cannot be used together with filtering (“include” and “exclude”) and/or monitoring (“watch”).

For more information on JSON usage, please refer to Section 4.18.2, “JSON API,” on page 577.

For a list of commands supporting the JSON API, please refer to Appendix B,“Show Commands Not Supported by JSON,” on page 1943.

3.2.8 CLI Shortcuts

Table 13 presents the available keyboard shortcuts on the Onyx CLI.

Table 13 - CLI Keyboard Shortcuts

Key Combination	Description
Ctrl-a	Move cursor to beginning of line
Ctrl-b	Move cursor backward one character without deleting
Ctrl-c	Terminate operation
Ctrl-d	If cursor is in the middle of the line, delete one character forward If cursor is at the end of the line, show auto-complete options for current word or word fragment If cursor at an empty line, same as Esc
Ctrl-e	Move cursor to end of line
Ctrl-f	Move cursor forward one character
Ctrl-h	Delete one character backwards from cursor
Ctrl-i	Auto-complete current word (same as TAB)
Ctrl-j	Return carriage (same as ENTER)
Ctrl-k	Delete line after cursor
Ctrl-l	Clear screen and show line at the top of terminal window
Ctrl-m	Return carriage (same as ENTER)
Ctrl-n	Next line (same as DOWN ARROW)
Ctrl-p	Next line (same as UP ARROW)
Ctrl-t	Transpose the two characters on either side of cursor
Ctrl-u	Delete line
Ctrl-y	Retrieve (“yank”) last item deleted
Esc b	Move cursor one word backward

Table 13 - CLI Keyboard Shortcuts

Key Combination	Description
Esc c	Capitalizes first letter in word after cursor
Esc d	Delete one word forward from cursor
Esc f	Move one word forward from cursor
Esc l	Change word after cursor to lowercase letters
Esc Ctrl-h	Delete one word backward from cursor
Esc [A	Next line (same as DOWN ARROW)
Esc [B	Next line (same as UP ARROW)
Esc [C	Move forward one character from cursor
Esc [D	Move backward one character from cursor

3.3 Web Interface Overview

The Onyx package equipped with web interface which is a web GUI that accept input and provide output by generating webpages which can be viewed by the user using a web browser.

The web interface makes available the following perspective tabs:

- Setup
- System
- Security
- Ports
- Status
- Ethernet Management
- IP Route

Make sure to save your changes before switching between menus or submenus. Click the “Save” button to the right of “Save Changes?”.

Figure 8: WebUI

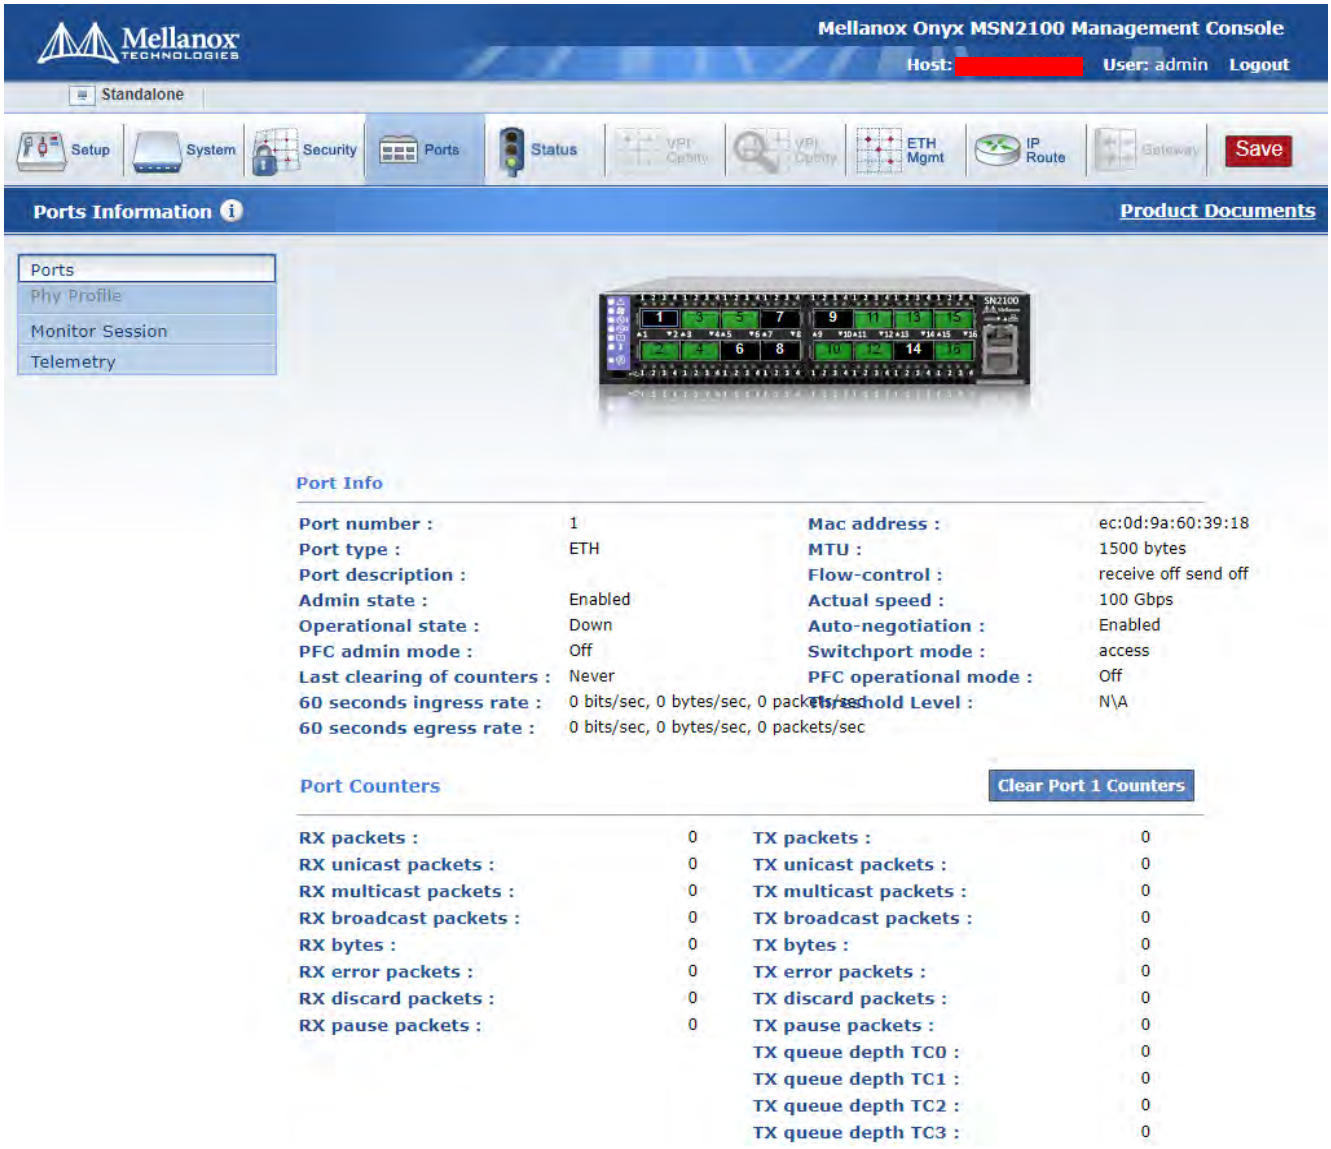

3.3.1 Setup Menu

The **Setup** menu makes available the following submenus (listed in order of appearance from top to bottom):

Table 14 - WebUI Setup Submenus

Submenu Title	Description
Interfaces	Obtains the status of, configures, or disables interfaces to the fabric. Thus, you can: set or clear the IP address and netmask of an interface; enable DHCP to dynamically assign the IP address and netmask; and set interface attributes such as MTU, speed, duplex, etc.

Table 14 - WebUI Setup Submenus

Submenu Title	Description
Routing	Configures, removes or displays the default gateway, and the static and dynamic routes
Hostname	Configures or modifies the hostname Configures or deletes static hosts **Note:** Changing hostname stamps a new HTTPS certificate
DNS	Configures, removes, modifies or displays static and dynamic name servers
Login Messages	Edits the login messages: Message of the Day (MOTD), Remote Login message, and Local Login message
Address Resolution	Adds static and dynamic ARP entries, and clears the dynamic ARP cache
IPSec	Configures IPSec
Neighbors	Displays IPv6 neighbor discovery protocol
Virtualization	Manages the virtualization and virtual machines
Virtual Switch Mgmt	Configures the system profile
Web	Configures web user interface and proxy settings
SNMP	Configures SNMP attributes, SNMP admin user, and trap sinks
Email Alerts	Configures the destination of email alerts and the recipients to be notified
XML gateway	Provides an XML request-response protocol to get and set hardware management information
JSON API	Manages JSON API
Logging	Sets up system log files, remote log sinks, and log formats
Configurations	Manages, activates, saves, and imports Onyx configuration files, and executes CLI commands
Docker	Manages docker images and containers.
Date and Time	Configures the date, time, and time zone of the switch system
NTP	Configures NTP (Network Time Protocol) and NTP servers
Licensing	Manages Onyx licenses

3.3.2 System Menu

The **System** menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Table 15 - WebUI System Submenus

Submenu Title	Description
Modules	Displays a graphic illustration of the system modules. By moving the mouse over the ports in the front view, a pop-up caption is displayed to indicate the status of the port. The port state (active/down) is differentiated by a color scheme (green for active, gray/black for down). By moving the mouse over the rear view, a pop-up caption is displayed to indicate the leaf part information.

Table 15 - WebUI System Submenus

Submenu Title	Description
Inventory	Displays a table with the following information about the system modules: module name, type, serial number, ordering part number and ASIC firmware version
Power Management	Displays a table with the following information about the system power supplies: power supply name, power, voltage level, current consumption, and status. A total power summary table is also displayed providing the power used, the power capacity, and the power available.
Onyx Upgrade	Displays the installed Onyx images (and the active partition), uploads a new image, and installs a new image
Reboot	Reboots the system. Make sure that you save your configuration prior to clicking reboot.

3.3.3 Security Menu

The **Security** menu makes available the following submenus (listed in order of appearance from top to bottom):

Table 16 - WebUI Security Submenus

Submenu Title	Description
Users	Manages (setting up, removing, modifying) user accounts
Admin Password	Modifies the system administrator password
SSH	Displays and generate host keys
AAA	Configures AAA (Authentication, Authorization, and Accounting) security services such as authentication methods and authorization
Login Attempts	Manages login attempts
RADIUS	Manages Radius client
TACACS+	Manages TACACS+ client
LDAP	Manages LDAP client
Certificate	Manages certificates

3.3.4 Ports Menu

The Ports menu displays the port state and enables some configuration attributes of a selected port. It also enables modification of the port configuration. A graphical display of traffic over time (last hour or last day) through the port is also available.

Table 17 - WebUI Ports Submenus

Submenu Title	Description
Ports	Manages port attributes, counters, transceiver info and displays a graphical counters histogram
Phy Profile	Provides the ability to manage PHY profiles

Table 17 - WebUI Ports Submenus

Submenu Title	Description
Monitor Session	Displays monitor session summary and enables configuration of a selected session
Telemetry	Displays and configures telemetry

3.3.5 Status Menu

The **Status** menu makes available the following submenus (listed in order of appearance from top to bottom):

Table 18 - WebUI Status Submenus

Submenu Title	Description
Summary	Displays general information about the switch system and the Onyx image, including current date and time, hostname, uptime of system, system memory, CPU load averages, etc.
Profile and Capabilities	Displays general information about the switch system capabilities such as the enabled profiles (e.g IB/ETH) and their corresponding values
What Just Happened	Displays and configures What Just Happened packet drop reasons.
Temperature	Provides a graphical display of the switch module sensors' temperature levels over time (1 hour). It is possible to display either the temperature level of one module's sensor or the temperature levels of all the module sensors' together.
Power Supplies	Provides a graphical display of one of the switch's power supplies voltage level over time (1 hour)
Fans	Provides a graphical display of fan speeds over time (1 hour). The display is per fan unit within a fan module.
CPU Load	Provides a graphical display of the management CPU load over time (1 hour)
Memory	Provides a graphical display of memory utilization over time (1 day)
Network	Provides a graphical display of network usage (transmitted and received packets) over time (1 day). It also provides per interface statistics.
Logs	Displays the system log messages. It is possible to display either the currently saved system log or a continuous system log.
Maintenance	Performs specific maintenance operations automatically on a predefined schedule
Alerts	Displays a list of the recent health alerts and enables the user to configure health settings
Virtualization	Displays the virtual machines, networks and volumes

3.3.6 ETH Mgmt

The **ETH Mgmt** menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Table 19 - WebUI ETH Mgmt Submenus

Submenu Title	Description
Spanning Tree	Configures and monitors spanning tree protocol
MAC Table	Configures static mac addresses in the switch, and displays the MAC address table
Link Aggregation	Configures and monitors aggregated Ethernet links (LAG) and configures LACP
VLAN	Manages the switch VLAN table
MLAG	Manages multi-chassis LAGs
IGMP Snooping	Manages IGMP snooping in the switch
ACL	Manages Access Control in the switch
Priority Flow Control	Manages priority flow control

3.3.7 IP Route

The **IP Route** menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Table 20 - WebUI IP Route Submenus

Submenu Title	Description
Router Global	Enables/disables IP routing protocol
IP Route	Configures, removes, and displays the routing table for router interfaces
IP Interface	Displays router interfaces
Address Resolution	Displays the address resolution (ARP) table for router interfaces
IP Diagnostic	Not implemented

3.4 Secure Shell (SSH)

It is recommended not to use more than 50 concurrent SSH sessions to the switch.

3.4.1 Adding a Host and Providing an SSH Key

➢ ***To add entries to the global known-hosts configuration file and its SSH value:***

Step 1. Change to Config mode Run:

```
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
switch [standalone: master] (config) #
```

Step 2. Add an entry to the global known-hosts configuration file and its SSH value. Run:

```
switch [standalone: master] (config) # ssh client global known-host "myserver ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAsXeklqc8T0EN2mnMcVcfhueaRYzIVqt4rVsrERIjmlJh4mkYYIa8hGGikN
a+t5xw2dRrNxnHYLK51bUsSG1ZNwZT1Dpme3pAZeMY7G4ZMgGIW9xOuaXgAA3eBeoUjFdi6+1BqchWk0nTb+gM
fI/MK/heQNns7AtTrvqg/O5ryIc="
switch [standalone: master] (config) #
```

Step 3. Verify what keys exist in the host. Run:

```
switch [standalone: master] (config) # show ssh client
SSH client Strict Hostkey Checking: ask

SSH Global Known Hosts:
    Entry 1: myserver
          Finger Print: d5:d7:be:d7:6c:b1:e4:16:df:61:25:2f:b1:53:a1:06

No SSH user identities configured.

No SSH authorized keys configured.

switch [standalone: master] (config) #
```

3.4.2 Retrieving Return Codes when Executing Remote Commands

➢ ***To stop the CLI and set the system to send return errors if some commands fail:***

Step 1. Connect to the system from the host SSH.

Step 2. Add the `-h` parameter after the `cli` (as shown in the example below) to notify the system to halt on failure and pass through the exit code.

```
ssh <username>@<hostname> cli -h '"enable" "show interfaces brief"'
```

3.5 Management Information Bases (MIBs)

The inventory in the switch system can be accessed through a MIB browser. These devices are indexed (entPhysicalIndex) using three levels:

1. Module layer which includes modules located on system (e.g. cables, fan, power supply, etc.). See table Table 21 for more details.
2. Device layer which includes system devices (e.g. switch devices, sensor aggregators, etc.). See table Table 22 for more details.
3. Sensor layer which includes system sensors (e.g. fan, and temperature sensors) located in the devices. See table Table 23 for more details.

Each layer is assigned a fixed position in the index number to represent it.

Figure 9: Index Scheme

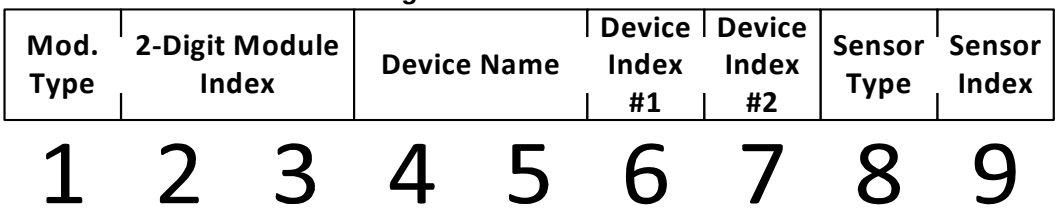

Each position could indicate different types of component according to the following criteria:

Table 21 - Module Type

Number	Description
1	Chassis
2	Management
3	Spine
4	Leaf
5	Fan
6	Power supply
7	BBU
8	x86 CPU
9	Port module

Table 22 - Device Type

Number	Description
01	PS
02	FAN
03	BOARD_MONITOR
04	CPU_BOARD_MONITOR
05	SX

Table 22 - Device Type

Number	Description
06	SIB
07	CPU_MEZZ_TEMP
08	CPU Package Sensor
09	CPU Core Sensor
10	SX_AMBIENT_TEMP
11	SX_MONITOR
12	AUX_IN_TMP_SNSR
13	AUX_OUT_TMP_SNSR
14	MAIN_IN_TMP_SNSR
15	MAIN_OUT_TMP_SNSR
16	CPU_MEZZ_TEMP
17	Controller
18	QSFP_TEMP
19	QSFP-ASIC
20	Board AMB temp
21	Ports AMB temp
22	Power monitor
23	PS_MONITOR
24	SWB AMB temp
25	pcie-switch-temp
26	SPC

Table 23 - Sensor Type

Number	Description
1	t – temperature sensor
2	f – fan sensor

For example:

- 401191311

  The first layer is “401” where:

  - “4”, according to Table 21, indicates a leaf
  - “01” indicates index #1 (Leaf #1)

  The second layer is “1913” where:

  - “19”, according to Table 22, indicates a QSFP ASIC
  - “1” indicates ASIC #1

- "3" indicates sensor #3 (QSFP-ASIC1-3)

The third layer is "11" where:

- "1", according to Table 23, indicates a temperature sensor
- "1" indicates sensor #1 (T1)

The resulting output in the entPhysicalDescr column of the MIB would be: L01/QSFP-ASIC-1/T1.

- 501020021

The first layer is 501 where

- "5", according to Table 21, indicates a fan
- "01 indicates index #1 (Fan #1)

The second layer is 0200 where:

- 02, according to Table 22, indicates a fan
- 0 – indicates that there is no first index
- 0 – indicates that there is no second index

The third layer is 21 where:

- "2", according to Table 23, indicates a fan sensor
- "1" indicates sensor #1 (F1)

The resulting output in the entPhysicalDescr column of the MIB would be: FAN1/FAN/F1.

3.6 Commands

3.6.1 CLI Session

This chapter displays all the relevant commands used to manage CLI session terminal.

cli clear-history

	cli clear-history Clears the command history of the current user.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # cli clear-history` `switch (config) #`
Related Commands	N/A
Note	

cli default

cli default {auto-logout <minutes> | paging enable | prefix-modes {enable | show-config} | progress enable | prompt {confirm-reload | confirm-reset | confirm-unsaved | empty-password}}
no cli default {auto-logout | paging enable | prefix-modes {enable | show-config} | progress enable prompt {confirm-reload | confirm-reset | confirm-unsaved | empty-password}

Configures default CLI options for all future sessions.
The no form of the command deletes or disables the default CLI options.

Syntax Description	minutes	Configures keyboard inactivity timeout for automatic logout. Range is 0-35791 minutes. Setting the value to 0 or using the no form of the command disables the auto-logout.
	paging enable	Enables text viewing one screen at a time.
	prefix-modes {enable | show-config}	Configures the prefix modes feature of CLI. • "prefix-modes enable" enables prefix modes for current and all future sessions • "prefix-modes show-config" uses prefix modes in "show configuration" output for current and all future sessions
	progress enable	Enables progress updates.
	prompt confirm-reload	Prompts for confirmation before rebooting.
	prompt confirm-reset	Prompts for confirmation before resetting to factory state.
	prompt confirm-unsaved	Confirms whether or not to save unsaved changes before rebooting.
	prompt empty-password	Prompts for a password if none is specified in a pseudo-URL for SCP.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # cli default prefix-modes enable`	
Related Commands	show cli	
Note		

cli max-sessions

cli max-sessions <number>
no cli max-sessions

Configures the maximum number of simultaneous CLI sessions allowed.
The no form of the command resets this value to its default.

Syntax Description	number	Range: 3-30
Default	30 sessions	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # cli max-sessions 40`	
Related Commands	show terminal	
Note		

cli session

cli session {auto-logout <minutes> | paging enable | prefix-modes {enable | show-config} | progress enable | terminal {length <size> | resize | type <terminal-type> | width} | x-display full <display>}
no cli session {auto-logout | paging enable | prefix-modes {enable | show-config} | progress enable | terminal type | x-display}

Configures default CLI options for all future sessions.
The no form of the command deletes or disables the CLI sessions.

Syntax Description	minutes	Configures keyboard inactivity timeout for automatic logout. Range is 0-35791 minutes. Setting the value to 0 or using the no form of the command disables the auto logout.
	paging enable	Enables text viewing one screen at a time.
	prefix-modes enable | show-config	Configures the prefix modes feature of CLI. • "prefix-modes enable" enables prefix modes for current and all future sessions • "prefix-modes show-config" uses prefix modes in "show configuration" output for current and all future sessions
	progress enable	Enables progress updates.
	terminal length	Sets the number of lines for the current terminal. Valid range is 5-999.
	terminal resize	Resizes the CLI terminal settings (to match the actual terminal window).
	terminal-type	Sets the terminal type. Valid options are: • ansi • console • dumb • linux • unknown • vt52 • vt100 • vt102 • vt220 • vt320 • xterm
	terminal width	Sets the width of the terminal in characters. Valid range is 34-999.
	x-display full <display>	Specifies the display as a raw string, e.g localhost:0.0.
Default	N/A	

Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # cli session auto-logout`
Related Commands	show terminal
Note	

terminal

terminal {length <number of lines> | resize | type <terminal type> | width <number of characters>}
no terminal type

Configures default CLI options for all future sessions.
The no form of the command clears the terminal type.

Syntax Description	length	Sets the number of lines for this terminal Range: 5-999
	resize	Resizes the CLI terminal settings (to match with real terminal)
	type	Sets the terminal type. Possible values: ansi, console, dumb, linux, screen, vt52, vt100, vt102, vt220, xterm.
	width	Sets the width of this terminal in characters Range: 34-999
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # terminal length 500` `switch (config) #`	
Related Commands	show terminal	
Note		

terminal sysrq enable

	terminal sysrq enable **no terminal sysrq enable** Enable SysRq over the serial connection (RS232 or Console port). The no form of the command disables SysRq over the serial connection (RS232 or Console port).
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.4.3000
Role	admin
Example	`switch (config) # terminal sysrq enable` `switch (config) #`
Related Commands	show terminal
Note	

show cli

	show cli Displays the CLI configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show cli`
	`CLI current session settings:`
	`  Maximum line size:  8192`
	`  Terminal width:  171 columns`
	`  Terminal length:  38 rows`
	`  Terminal type:  xterm`
	`  X display setting:  (none)`
	`  Auto-logout:  disabled`
	`  Paging:  enabled`
	`  Progress tracking:  enabled`
	`  Prefix modes:  disabled`
	
	`CLI defaults for future sessions:`
	`  Auto-logout:  disabled`
	`  Paging:  enabled`
	`  Progress tracking:  enabled`
	`  Prefix modes:  enabled (and use in 'show configuration')`
	
	`Settings for both this session and future ones:`
	`  Show hidden config:  yes`
	`  Confirm losing changes:  yes`
	`  Confirm reboot/shutdown:  no`
	`  Confirm factory reset:  yes`
	`  Prompt on empty password:  yes`
	`switch (config) #`

Related Commands	cli default
Note	

show cli max-sessions

	show cli max-sessions Displays maximum number of sessions.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200
Role	admin

Example	`switch (config) # show cli max-sessions`
	`Maximum number of CLI sessions: 5`
	`switch (config) #`

Related Commands	
Note	

show cli num-sessions

	show cli num-sessions Displays current number of sessions.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200
Role	admin

Example	`switch (config) # show cli num-sessions`
	`Current number of CLI sessions: 40`
	`switch (config) #`

Related Commands	
Note	

3.6.2 Banner

banner login

banner login <string>
no banner login

Sets the CLI welcome banner message.
The no form of the command resets the system login banner to its default.

Syntax Description	string Text string.
Default	"Mellanox Onyx Switch Management"
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # banner login Example`
Related Commands	show banner
Note	If more than one word is used (there is a space) quotation marks should be added (i.e. "xxxx xxxx").

banner login-local

banner login-local <string>
no banner login-local

Sets system login local banner.
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.5.0200	Added no form of the command
Role	admin	
Example	`switch (config) # banner login-local Testing` `switch (config) #`	
Related Commands	show banner	
Note	• The login-local refers to the serial connection banner • If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

banner login-remote

banner login-remote <string>
no banner login-remote

Sets system login remote banner.
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.5.0200	Added no form of the command
Role	admin	
Example	`switch (config) # banner login-remote Testing` `switch (config) #`	
Related Commands	show banner	
Note	• The login-remote refers to the SSH connections banner • If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

banner logout

banner logout <string>
no banner logout

Set system logout banner (for both local and remote logins).
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # banner logout Testing` `switch (config) #`	
Related Commands	show banner	
Note	If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

banner logout-local

banner logout-local <string>
no banner logout-local

Sets system logout local banner.
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # banner logout-local Testing` `switch (config) #`	
Related Commands	show banner	
Note	• The logout-local refers to the serial connection banner • If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

banner logout-remote

banner logout-remote <string>
no banner logout-remote

Sets system logout remote banner.
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # banner logout-remote Testing` `switch (config) #`	
Related Commands	show banner	
Note	• The logout-remote refers to SSH connections banner • If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

banner motd

banner motd <string>
no banner motd

Configures the message of the day banner.
The no form of the command resets the system Message of the Day banner.

Syntax Description	string	Text string
Default	“Mellanox Switch”	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # banner motd “My Banner”`	
Related Commands	show banner	
Note	• If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”). • To insert a multi-line MotD, hit Ctrl-V (escape sequence) followed by Ctrl-J (new line sequence). The symbol “^J” should appear. Then, whatever is typed after it becomes the new line of the MotD. Remember to also include the string between quotation marks.	

show banner

	show banner Displays configured banners.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.5.0200 Updated Example 3.6.6000 Updated Example
Role	Any command mode

Example	`switch (config) # show banner`
	
	`Banners:`
	`  Message of the Day (MOTD):`
	`  Mellanox Switch`
	
	`  Login:`
	`  Mellanox ONYX Switch Management`
	
	`  Logout:`
	`  Goodbye`

Related Commands	banner login banner login-local banner login-remote banner logout banner logout-local banner logout-remote banner motd
Note	

3.6.3 SSH

ssh server enable

ssh server enable
no ssh server enable

Enables the SSH server.
The no form of the command disables the SSH server.

Syntax Description	N/A
Default	SSH server is enabled
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # ssh server enable`
Related Commands	show ssh server
Note	Disabling SSH server does not terminate existing SSH sessions, it only prevents new ones from being established.

ssh server host-key

	ssh server host-key {<key-type> {private-key <private-key>| public-key <public-key>} | generate} Configures host keys for SSH.	
Syntax Description	key-type	• rsa1 - RSAv1 • rsa2 - RSAv2 • dsa2 - DSAv2
	private-key	Sets new private-key for the host keys of the specified type
	public-key	Sets new public-key for the host keys of the specified type
	generate	Generates new RSA and DSA host keys for SSH
Default	SSH keys are locally generated	
Configuration Mode	config	
History	3.1.0000	
	3.4.2300	Added notes
Role	admin	

Example	`switch (config) # ssh server host-key dsa2 private-key`
	`Key: ********************************************`
	`Confirm: ******************************************** `	

Related Commands	show ssh server system secure-mode enable	
Note	When working in secure mode, the commands “ssh server host-key rsa1” and “ssh server host-key generate” do not create RSAv1 key-type.	

ssh server login attempts

ssh server login attempts <number>
no ssh server login attempts

Configures maximum login attempts on SSH server.
The no form of the command resets the login attempts value to its default.

Syntax Description	number	Range: 3-100 attempts.
Default	6 attempts	
Configuration Mode	config	
History	3.5.0200	
	3.5.1000	Increased minimum number of attempts allowed
Role	admin	
Example	`switch (config) # ssh server login attempts 5`	
Related Commands	show ssh server	
Note		

ssh server login timeout

ssh server login timeout <time>
no ssh server login timeout

Configures login timeout on SSH server.
The no form of the command resets the timeout value to its default.

Syntax Description	time	Range: 1-600 seconds
Default	120 seconds	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # ssh server login timeout 130`	
Related Commands	show ssh server	
Note		

ssh server min-version

ssh server min-version <version>
no ssh server min-version

Sets the minimum version of the SSH protocol that the server supports.
The no form of the command resets the minimum version of SSH protocol supported.

Syntax Description	version	Possible versions are 1 and 2.
Default	2	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ssh server min-version 2`	
Related Commands	show ssh server	
Note		

ssh server ports

ssh server ports {<port1> [<port2>...]}

Specifies which ports the SSH server listens on.

Syntax Description	port	Port number in [1...65535].
Default	22	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ssh server ports 22`	
Related Commands	show ssh server	
Note	• Multiple ports can be specified by repeating the <port> parameter • The command will remove any previous ports if not listed in the command	

ssh server security strict

	ssh server security strict Enables strict security settings. The no form of the command disables strict security settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.5060 3.6.4000
Role	admin
Example	`switch (config) # ssh server security strict`
Related Commands	show ssh server
Note	The following ciphers are disabled for SSH when strict security is enabled: • aes256-cbc • aes192-cbc • aes128-cbc • arcfour • blowfish-cbc • cast128-cbc • rijndael-cbc@lysator.liu.se • 3des-cbc

ssh server tcp-forwarding enable

	ssh server tcp-forwarding enable Enables TCP port forwarding. The no form of the command disables TCP port forwarding.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # ssh server tcp-forwarding enable`
Related Commands	show ssh server
Note	

ssh server x11-forwarding

	ssh server x11-forwarding enable **no ssh server x11-forwarding enable** Enables X11 forwarding on the SSH server. The no form of the command disables X11 forwarding.
Syntax Description	N/A
Default	X11-forwarding is disabled.
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # ssh server x11-forwarding enable`
Related Commands	N/A
Note	

ssh client global

ssh client global {host-key-check <policy>} | known-host <known-host-entry>}
no ssh client global {host-key-check | known-host localhost}

Configures global SSH client settings.
The no form of the command negates global SSH client settings.

Syntax Description	host-key-check <policy>	Sets SSH client configuration to control how host key checking is performed. This parameter may be set in 3 ways. • If set to “no” it always permits connection, and accepts any new or changed host keys without checking • If set to “ask” it prompts user to accept new host keys, but does not permit a connection if there was already a known host entry that does not match the one presented by the host • If set to “yes” it only permits connection if a matching host key is already in the known hosts file
	known-host	Adds an entry to the global known-hosts configuration file.
	known-host-entry	Adds/removes an entry to/from the global known-hosts configuration file. The entry consist of “<IP> <key-type> <key>”.
Default	host-key-check - ask, no keys are configured by default	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # ssh client global host-key-check no
switch (config) # ssh client global known-host "72.30.2.2 ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEArB9i5OnukAHNUOkwpCmEl0m88kJgB-
zL22+F5tfaSn+S0pVYxrceZeyuzXsoZ1VtFTk2Fydwy0YvMS0Kcv2PuCrPZV/
GYd31QEnn22rEmrlPrKCrMl1XlUy6DFlr3OgwWm1baobmDlG/gSziWz/gc4Jgqf2CyX-
Fq4pzaR1jar1Vk="

switch (config) # show ssh client
SSH client Strict Hostkey Checking: ask

SSH Global Known Hosts:
    Entry 1: 72.30.2.2
          Finger Print: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6

No SSH user identities configured.

No SSH authorized keys configured.

switch (config) #
```

Related Commands	show ssh client
Note	

ssh client user

ssh client user <username> {authorized-key sshv2 <public key> | identity <key type> {generate | private-key [<private key>] | public-key [<public key>]} | known-host <known host> remove}
no ssh client user admin {authorized-key sshv2 <public key ID> | identity <key type>}

Adds an entry to the global known-hosts configuration file, either by generating new key, or by adding manually a public or private key.
The no form of the command removes a public key from the specified user's authorized key list, or changes the key type.

Syntax Description	username	The specified user must be a valid account on the system. Possible values for this parameter are “admin”, “monitor”, “xmladmin”, and “xmluser”.
	authorized-key sshv2 <public key>	Adds the specified key to the list of authorized SSHv2 RSA or DSA public keys for this user account. These keys can be used to log into the user's account.
	identity <key type>	Sets certain SSH client identity settings for a user, dsa2 or rsa2.
	generate	Generates SSH client identity keys for specified user.
	private-key	Sets private key SSH client identity settings for the user.
	public-key	Sets public key SSH client identity settings for the user.
	known-host <known host> remove	Removes host from user's known host file.
Default	No keys are created by default	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ssh client user admin known-host 172.30.1.116 remove`	
Related Commands	show ssh client	
Note	If a key is being pasted from a cut buffer and was displayed with a paging program, it is likely that newline characters have been inserted, even if the output was not long enough to require paging. One can specify “no cli session paging enable” before running the “show” command to prevent the newlines from being inserted.	

slogin

slogin [<slogin options>] <hostname>

Invokes the SSH client. The user is returned to the CLI when SSH finishes.

Syntax Description	slogin options	usage: slogin [-1246AaCfgkNnqsTtVvXxY] [-b bind_address] [-c cipher_spec] [-D port] [-e escape_char] [-F configfile] [-i identity_file] [-L port:host:hostport] [-l login_name] [-m mac_spec] [-o option] [-p port] [-R port:host:hostport] [user@]host-name [command]
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	monitor/admin	

Example	`switch (config) # slogin 192.168.10.70`
	`The authenticity of host '192.168.10.70 (192.168.10.70)' can't be estab-lished.`
	`RSA key fingerprint is 2e:ad:2d:23:45:4e:47:e0:2c:ae:8c:34:f0:1a:88:cb.`
	`Are you sure you want to continue connecting (yes/no)? yes`
	`Warning: Permanently added '192.168.10.70' (RSA) to the list of known hosts.`
	
	`Mellanox Onyx Switch Management`
	
	`Last login: Sat Feb 28 22:55:17 2009 from 10.208.0.121`
	
	`Mellanox Switch`
	
	`switch (config) #`	

Related Commands	N/A	
Note		

show ssh client

	show ssh client Displays the client configuration of the SSH server.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # show ssh client`
	`SSH client Strict Hostkey Checking: ask`
	
	`SSH Global Known Hosts:`
	`  Entry 1: 72.30.2.2`
	`  Finger Print: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6`
	
	`No SSH user identities configured.`
	
	`No SSH authorized keys configured.`
	
	`switch (config) #`

Related Commands	N/A
Note	

show ssh server

	show ssh server Displays SSH server configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
	3.4.0000 Updated Example
	3.5.0200 Added SSH login timeout and max attempts
	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ssh server`
	`SSH server configuration:`
	`SSH server enabled: yes`
	`Server security strict mode: no`
	`Minimum protocol version: 2`
	`TCP forwarding enabled: yes`
	`X11 forwarding enabled: no`
	`SSH login timeout: 120`
	`SSH login max attempts: 6`
	`SSH server ports: 22`
	
	`Interface listen enabled: yes`
	`Listen Interfaces:`
	`No interface configured.`
	
	`Host Key Finger Prints and Key Lengths:`
	`RSA v1 host key: SHA256:sMgangJjG9FmSch/9Y9aZ/WJ2wKf3c+SeF8XKgYYdCA (2048)`
	`RSA v2 host key: SHA256:gVu6qLW1ZifEp8wRer2jkvILZMGNl6VCYU3HqC1INC8 (2048)`
	`DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024)`

Related Commands	ssh server
Note	

show ssh server host-keys

	show ssh server host-keys Displays SSH host key configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ssh server host-keys`
	`SSH server configuration:`
	`SSH server enabled: yes`
	`Server security strict mode: no`
	`Minimum protocol version: 2`
	`TCP forwarding enabled: yes`
	`X11 forwarding enabled: no`
	`SSH login timeout: 120`
	`SSH login max attempts: 6`
	`SSH server ports: 22`
	
	`Interface listen enabled: yes`
	`Listen Interfaces:`
	`No interface configured.`
	
	`Host Key Finger Prints and Key Lengths:`
	`RSA v1 host key: SHA256:sMgangJjG9FmSch/9Y9aZ/WJ2wKf3c+SeF8XKgYYdCA (2048)`
	`RSA v2 host key: SHA256:gVu6qLW1ZifEp8wRer2jkvILZMGNl6VCYU3HqC1INC8 (2048)`
	`DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024)`
	
	`Host Keys:`
	`RSA v1 host key: "kebo-2100-1 2048 65537 21801469875<...>27851"`
	`RSA v2 host key: "kebo-2100-1 ssh-rsa AAAAB3Nza<...>KE5"`
	`DSA v2 host key: "kebo-2100-1 ssh-dss AAAAB3Nza<...>/s="`

Related Commands	ssh server host-keys
Note	

3.6.4 Remote Login

telnet

telnet

Logs into another system using telnet.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # telnet` `telnet>`
Related Commands	telnet-server
Note	

telnet-server enable

	telnet-server enable **no telnet-server enable** Enables the telnet server. The no form of the command disables the telnet server.
Syntax Description	N/A
Default	Telnet server is disabled
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # telnet-server enable`
	`switch (config) # show telnet-server`
	`Telnet server enabled: yes`

Related Commands	show telnet-server
Note	

show telnet-server

	show telnet-server Displays telnet server settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # show telnet-server`
	`Telnet server enabled: yes`
	`switch (config) #`

Related Commands	telnet-server enable
Note	

3.6.5 Web Interface

web auto-logout

web auto-logout <number of minutes>
no web auto-logout <number of minutes>

Configures length of user inactivity before auto-logout of a web session.
The no form of the command disables the web auto-logout (web sessions will never logged out due to inactivity).

Syntax Description	number of minutes	The length of user inactivity in minutes. 0 will disable the inactivity timer (same as a "no web auto-logout" command).
Default	60 minutes	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web auto-logout 60`	
Related Commands	show web	
Note	The no form of the command does not automatically log users out due to inactivity.	

web cache-enable

	web cache-enable **no web cache-enable** Enables web clients to cache webpages. The no form of the command disables web clients from caching webpages.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.4.1100
Role	admin
Example	`switch (config) # no web cache-enable`
Related Commands	N/A
Note	

web client cert-verify

	web client cert-verify **no web client cert-verify** Enables verification of server certificates during HTTPS file transfers. The no form of the command disables verification of server certificates during HTTPS file transfers.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # web client cert-verify`
Related Commands	N/A
Note	

web client ca-list

web client ca-list {<ca-list-name> | default-ca-list | none}
no web client ca-list

Configures supplemental CA certificates for verification of server certificates during HTTPS file transfers.
The no form of the command uses no supplemental certificates.

Syntax Description	ca-list-name	Specifies CA list to configure.
	default-ca-list	Configures default supplemental CA certificate list.
	none	Uses no supplemental certificates.
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # web client ca-list default-ca-list`	
Related Commands	N/A	
Note		

web enable

	web enable **no web enable** Enables the web-based management console. The no form of the command disables the web-based management console.
Syntax Description	N/A
Default	enable
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # web enable`
Related Commands	show web
Note	

web http

web http {enable | port <port number> | redirect}
no web http {enable | port | redirect}

Configures HTTP access to the web-based management console.
The no form of the command negates HTTP settings for the web-based management console.

Syntax Description	enable	Enables HTTP access to the web-based management console.
	port number	Sets a port for HTTP access.
	redirect	Enables redirection to HTTPS. If HTTP access is enabled, this specifies whether a redirect from the HTTP port to the HTTPS port should be issued to mandate secure HTTPS access.
Default	HTTP is disabled HTTP TCP port is 80 HTTP redirect to HTTPS is disabled	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web http enable`	
Related Commands	show web web enable	
Note	Enabling HTTP is meaningful if the WebUI as a whole is enabled.	

web httpd

web httpd listen {enable | interface <ifName>}
no web httpd listen {enable | interface <ifName>}

Enables the listen interface restricted list for HTTP and HTTPS.
The no form of the command disables the HTTP server listen ability.

Syntax Description	enable	Enables Web interface restrictions on access to this system.
	interface <ifName>	Adds interface to Web server access restriction list (i.e. mgmt0, mgmt1)
Default	Listening is enabled. all interfaces are permitted.	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web httpd listen enable`	
Related Commands	N/A	
Note	If enabled, and if at least one of the interfaces listed is eligible to be a listen interface, then HTTP/HTTPS requests will only be accepted on those interfaces. Otherwise, HTTP/HTTPS requests are accepted on any interface.	

web https

web https {certificate {regenerate | name | default-cert} | enable | port <port number> | ssl ciphers {all | TLS | TLS1.2}}
no web https {enable | port <port number>}

Configures HTTPS access to the web-based management console.
The no form of the command negates HTTPS settings for the web-based management console.

Syntax Description	certificate regenerate	Re-generates certificate to use for HTTPS connections.
	certificate name	Configure the named certificate to be used for HTTPS connections
	certificate default-cert	Configure HTTPS to use the configured default certificate
	enable	Enables HTTPS access to the web-based management console.
	port	Sets a TCP port for HTTPS access.
	ssl ciphers {all | TLS | TLS1.2}	Sets ciphers to be used for HTTPS.
Default	HTTPS is enabled Default port is 443	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added “ssl ciphers” parameter
	3.4.0010	Added TLS parameter to “ssl ciphers”
Role	admin	
Example	`switch (config) # web https enable`	
Related Commands	show web web enable	
Note	• Enabling HTTPS is meaningful if the WebUI as a whole is enabled. • See the command “crypto certificate default-cert name” for how to change the default certificate if inheriting the configured default certificate is preferred	

web https ssl renegotiation enable

	web https ssl renegotiation enable **no web https ssl renegotiation enable** Enables SSL renegotiation flag in httpd web server. The no form of the command disables SSL renegotiation flag in httpd web server.
Syntax Description	N/A
Default	HTTPS is enabled Default port is 443
Configuration Mode	config
History	3.6.8008
Role	admin
Example	`switch (config) # web https ssl renegotiation enable`
Related Commands	show web web enable
Note	

web https ssl secure-cookie enable

	web https ssl secure-cookie enable **no web https ssl secure-cookie enable** Enables SSL secure-cookie flag in httpd web server. The no form of the command disables secure-cookie flag in httpd web server.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.8008
Role	admin
Example	`switch (config) # web https ssl secure-cookie enable`
Related Commands	show web web enable
Note	

web proxy auth authtype

web proxy auth authtype <auth-type>
no web proxy auth authtype

Configures type of authentication to use with web proxy.
The no form of the command resets web proxy authentication type to its default.

Syntax Description	auth-type	Possible values: • none – no authentication • basic – HTTP basic authentication
Default	Basic authentication settings	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web proxy auth authtype basic`	
Related Commands	show web web enable	
Note		

web proxy auth basic

web proxy auth basic {password <password> | username <username>}
no web proxy auth basic {password | username}

Configures HTTP basic authentication settings for proxy.
The no form of the command clears password or username configuration.

Syntax Description	password	Sets plaintext password for HTTP basic authentication with web proxy
	username	Sets username for HTTP basic authentication with web proxy
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web proxy auth basic password 57R0ngP455w0rD`	
Related Commands	show web web enable	
Note		

web proxy auth host

web proxy auth host <ip-address> [port <number>]

Configures web proxy auth host.

Syntax Description	port	Sets web proxy default port
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # web proxy auth host 2001:0db8:85a3::8a2e:0370:7334 port 3`	
Related Commands	show web web enable	
Note		

show web

	show web Displays WebUI configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000 3.6.8008 Updated Example
Role	admin

Example	`switch (config) # show web`
	`Web User Interface:`
	`  Web interface enabled:  yes`
	`  Web caching enabled:  no`
	`  HTTP enabled:  no`
	`  HTTP port:  80`
	`  HTTP redirect to HTTPS:  no`
	`  HTTPS enabled:  yes`
	`  HTTPS port:  443`
	`  HTTPS ssl-ciphers:  TLS1.2`
	`  HTTPS ssl-renegotiation: no`
	`  HTTPS ssl-secure-cookie: yes`
	`  HTTPS certificate name:  default-cert`
	`  Listen enabled:  yes`
	`  Listen Interfaces:`
	`  No interface configured.`
	
	`  Inactivity timeout:  1 hr`
	`  Session timeout:  2 hr 30 min`
	`  Session renewal:  30 min`
	
	`Web file transfer proxy:`
	`  Proxy enabled: no`
	
	`Web file transfer certificate authority:`
	`  HTTPS server cert verify: yes`
	`  HTTPS supplemental CA list: default-ca-list`

Related Commands	
Note	

4 System Management

4.1 Management Interface

Management interfaces are used in order to provide access to switch management user interfaces (e.g. CLI, WebUI). Mellanox switches support out-of-band (OOB) dedicated interfaces (e.g. mgmt0, mgmt1) and in-band dedicated interfaces. In addition, most Mellanox switches feature a serial port that provides access to the CLI only.

On switch systems with two OOB management ports, both of them may be configured on the same VLAN if needed. In this case, ARP replies to the IP of those management interfaces is answered from either of them.

4.1.1 Configuring Management Interfaces with Static IP Addresses

If your switch system was set during initialization to obtain dynamic IP addresses through DHCP and you wish to switch to static assignments, perform the following steps:

Step 1. Enter Config configuration mode. Run:

```
switch >
switch > enable
switch # configure terminal
switch (config) #
```

Step 2. Disable setting IP addresses using the DHCP using the following command:

```
switch (config) # no interface <ifname> dhcp
```

Step 3. Define your interfaces statically using the following command:

```
switch (config) # interface <ifname> ip address <IP address> <netmask>
```

4.1.2 Configuring IPv6 Address on the Management Interface

Step 1. Enable IPv6 on this interface. Run:

```
switch (config) # interface mgmt0 ipv6 enable
```

Step 2. Set the IPv6 address to be configured automatically. Run:

```
switch (config) # interface mgmt0 ipv6 address autoconfig
```

Step 3. Verify the IPv6 address is configured correctly. Run:

```
switch (config) # show interfaces mgmt0 brief
```

4.1.3 Dynamic Host Configuration Protocol (DHCP)

DHCP is used for automatic retrieval of management IP addresses.

For all other systems (and software versions) DHCP is disabled by default.

If a user connects through SSH, runs the wizard and turns off DHCP, the connection is immediately terminated as the management interface loses its IP address.

```
<localhost># ssh admin@<ip-address>
Mellanox Onyx Switch Management
Password:
Mellanox switch
Mellanox configuration wizard
Do you want to use the wizard for initial configuration? yes
Step 1: Hostname? [my-switch]
Step 2: Use DHCP on mgmt0 interface? [yes] no
<localhost>#
```

In such case the serial connection should be used.

4.1.4 Default Gateway

To configure manually the default gateway, use the “ip route” command, with “0.0.0.0” as prefix and mask. The next-hop address must be within the range of one of the IP interfaces on the system.

```
switch (config)# ip route 0.0.0.0 0.0.0.0 10.10.0.2
switch (config)# show ip route
Destination      Mask             Gateway         Interface   Source   Distance/Metric
default          0.0.0.0           10.10.0.2      mgmt0       static   0/0
10.10.0.0        255.255.254.0     0.0.0.0        mgmt0       direct   0/0
switch (config)#
```

4.1.5 In-Band Management

In-band management is a management path passing through the data ports. In-band management can be created over one of the VLANs in the systems.

The in-band management feature does not require any license. However, it works only for the system profile Ethernet. It can be enabled with IP Routing.

➢ ***To set an in-band management channel:***

Step 1. Create a VLAN. Run:

```
switch (config)# vlan 10
switch (config vlan 10) #
```

Step 2. Create a VLAN interface. Run:

```
switch (config) # interface vlan 10
switch (config interface vlan 10) #
```

Step 3. Configure L3 attributes on the newly created VLAN interface. Run:

```
switch (config interface vlan 10) # ip address 10.10.10.10 /24
```

Step 4. (Optional) Verify in-band management configuration. Run:

```
switch (config) # show interfaces vlan 10
 Admin state: Enabled
 Operational state: Up
 Mac Address: f4:52:14:67:07:e8
 Internet Address: 10.10.10.10/24
 Broadcast address: 10.10.10.255
 MTU: 1500 bytes
 Arp timeout: 1500 seconds
 Icmp redirect: Disabled
 Description: N/A
 VRF: default
 Counters: Enabled
RX
 0 Unicast packets
 0 Multicast packets
 0 Unicast bytes
 0 Multicast bytes
 0 Bad packets
 0 Bad bytes
TX
 0 Unicast packets
 0 Multicast packets
 0 Unicast bytes
 0 Multicast bytes
switch (config) #
```

4.1.6 Configuring Hostname via DHCP (DHCP Client Option 12)

This feature, also known as the DHCP Client Option 12, is enabled by default and assigns the switch system a hostname via DHCP as long as network manager configures hostname to the management interfaces' (i.e. mgmt0, mgmt1) MAC address. If a network manager configures the hostname manually through any of the user interfaces, the hostname is not retrieved from the DHCP server.

- ***To enable fetching hostname from DHCP server, run:***

```
switch (config interface mgmt0) # dhcp hostname
```

- ***To disable fetching hostname from DHCP server, run:***

```
switch (config interface mgmt0) # no dhcp hostname
```

Getting the hostname through DHCP is enable by default and will change the switch hostname if the hostname is not set by the user. Therefore, if a switch is part of an HA cluster the user would need to make sure the HA master has the same HA node names as the DHCP server.

4.1.7 Commands

4.1.7.1 Interface

This chapter describes the commands should be used to configure and monitor the management interface.

interface

interface {mgmt0 | mgmt1 | lo | vlan<id>}

Enters a management interface context.

Syntax Description	mgmt0	Management port 0 (out of band).
	mgmt1	Management port 1 (out of band).
	lo	Loopback interface.
	vlan<id>	In-band management interface (e.g. vlan10).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # interface mgmt0` `switch (config interface mgmt0) #`	
Related Commands	show interfaces <ifname>	
Notes		

ip address

ip address <IP address> <netmask>
no ip address

Sets the IP address and netmask of this interface.
The no form of the command clears the IP address and netmask of this interface.

Syntax Description	IP address	IPv4 address
	netmask	Subnet mask of IP address
Default	0.0.0.0/0	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config) # interface mgmt0` `switch (config interface mgmt0) # ip address 10.10.10.10 255.255.255.0`	
Related Commands	show interfaces <ifname>	
Notes	If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled.	

ip default-gateway

	ip default-gateway <next hop IP address or interface name> **no ip default-gateway** Configures a default route. The no form of the command removes the current default route.
Syntax Description	next hop IP address or interface name — IP address, lo, mgmt0, or mgmt1.
Default	N/A
Configuration Mode	config interface management
History	3.1.0000
Role	admin
Example	`switch (config) # ip default-gateway mgmt1` `switch (config) #`
Related Commands	
Notes	

alias

alias <index> ip address < IP address> <netmask>
no alias <index>

Adds an additional IP address to the specified interface. The secondary address will appear in the output of "show interface" under the data of the primary interface along with the alias.
The no form of the command removes the secondary address to the specified interface.

Syntax Description	index	A number that is to be aliased to (associated with) the secondary IP.
	IP address	Additional IP address.
	netmask	Subnet mask of the IP address.
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # alias 2 ip address 9.9.9.9 255.255.255.255`	
Related Commands	show interfaces <ifname>	
Notes	• If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled • More than one additional IP address can be added to the interface	

mtu

mtu <bytes>
no mtu <bytes>

Sets the Maximum Transmission Unit (MTU) of this interface.
The no form of the command resets the MTU to its default.

Syntax Description	bytes	The entry range is 68-1500.
Default	1500	
Configuration Mode	config interface management	
History	3.6.3004	
Role	admin	
Example	`switch (config interface mgmt0) # mtu 1500`	
Related Commands	show interfaces <ifname>	
Notes		

duplex

duplex <duplex>
no duplex

Sets the interface duplex.
The no form of the command resets the duplex setting for this interface to its default value.

Syntax Description	duplex	Sets the duplex mode of the interface. The following are the possible values: • half - half duplex • full - full duplex • auto - auto duplex sensing (half or full)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # duplex auto`	
Related Commands	show interfaces <ifname>	
Notes	• Setting the duplex to “auto” also sets the speed to “auto” • Setting the duplex to one of the settings “half” or “full” also sets the speed to a manual setting which is determined by querying the interface to find out its current auto-detected state	

speed

speed <speed>
no speed

Sets the interface speed.
The no form of the command resets the speed setting for this interface to its default value.

Syntax Description	speed	Sets the speed of the interface. The following are the possible values: • 10 - fixed to 10Mbps • 100 - fixed to 1000Mbps • 1000 - fixed to 1000Mbps • auto - auto speed sensing (10/100/1000Mbps)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # speed auto`	
Related Commands	show interfaces <ifname>	
Notes	• Setting the speed to “auto” also sets the duplex to “auto” • Setting the speed to one of the manual settings (generally “10”, “100”, or “1000”) also sets the duplex to a manual setting which is determined by querying the interface to find out its current auto-detected state	

dhcp

dhcp [renew]
no dhcp

Enables DHCP on the specified interface.
The no form of the command disables DHCP on the specified interface.

Syntax Description	renew	Forces a renewal of the IP address. A restart on the DHCP client for the specified interface will be issued.
Default	Could be enabled or disabled (per part number) manufactured with 3.2.0500	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # dhcp`	
Related Commands	show interfaces <ifname> configured	
Notes	• When enabling DHCP, the IP address and netmask are received via DHCP hence, the static IP address configuration is ignored • Enabling DHCP disables zeroconf and vice versa • Setting a static IP address and netmask does not disable DHCP. DHCP is disabled using the “no” form of this command, or by enabling zeroconf.	

dhcp hostname

	dhcp hostname **no dhcp hostname** Enables fetching the hostname from DHCP for this interface. The no form of the command disables fetching the hostname from DHCP for this interface.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface management
History	3.5.1000
Role	admin
Example	`switch (config interface mgmt0) # dhcp hostname` `switch (config interface mgmt0) #`
Related Commands	hostname <hostname> show interfaces <ifname> configured
Notes	• If a hostname is configured manually by the user, that configuration would override the “dhcp hostname” configuration • After upgrading to version 3.5.1000 when a default hostname is not configured, the DHCP server assigns the new hostname for your machine • These commands do not work on in-band interfaces

shutdown

	shutdown **no shutdown** Disables the specified interface. The no form of the command enables the specified interface.
Syntax Description	N/A
Default	no shutdown
Configuration Mode	config interface management
History	3.1.0000
Role	admin
Example	`switch (config interface mgmt0) # no shutdown`
Related Commands	show interfaces <ifname> configured
Notes	

zeroconf

	zeroconf **no zeroconf** Enables zeroconf on the specified interface. It randomly chooses a unique link-local IPv4 address from the 169.254.0.0/16 block. This command is an alternative to DHCP. The no form of the command disables the use of zeroconf on the specified interface.
Syntax Description	N/A
Default	no zeroconf
Configuration Mode	config interface management
History	3.1.0000
Role	admin
Example	`switch (config interface mgmt0) # zeroconf`
Related Commands	show interfaces <ifname> configured
Notes	Enabling zeroconf disables DHCP and vice versa.

comment

comment <comment>
no comment

Adds a comment for an interface.
The no form of the command removes a comment for an interface.

Syntax Description	comment	A free-form string that has no semantics other than being displayed when the interface records are listed.
Default	no comment	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # comment my-interface`	
Related Commands	N/A	
Notes		

ipv6 enable

	ipv6 enable **no ipv6 enable** Enables all IPv6 addressing for this interface. The no form of the command disables all IPv6 addressing for this interface.
Syntax Description	N/A
Default	IPv6 addressing is disabled
Configuration Mode	config interface management
History	3.1.0000
Role	admin
Example	`switch (config interface mgmt0) # ipv6 enable`
Related Commands	ipv6 address show interface <ifname>
Notes	• The interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface • If IPv6 is enabled on an interface, the system will automatically add a link-local address to the interface. Link-local addresses can only be used to communicate with other hosts on the same link, and packets with link-local addresses are never forwarded by a router. • A link-local address, which may not be removed, is required for proper IPv6 operation. The link-local addresses start with “fe80::”, and are combined with the interface identifier to form the complete address.

ipv6 address

ipv6 address {<IPv6 address/netmask> | autoconfig [default | privacy]}
no ipv6 {<IPv6 address/netmask> | autoconfig [default | privacy]}

Configures IPv6 address and netmask to this interface, static or autoconfig options are possible.
The no form of the command removes the given IPv6 address and netmask or disables the autoconfig options.

Syntax Description	IPv6 address/netmask	Configures a static IPv6 address and netmask. Format example: 2001:db8:1234::5678/64.
	autoconfig	Enables IPv6 stateless address auto configuration (SLAAC) for this interface. An address will be automatically added to the interface based on an IPv6 prefix learned from router advertisements, combined with an interface identifier.
	autoconfig default	Enables default learning routes. The default route will be discovered automatically, if the autoconfig is enabled.
	autoconfig privacy	Uses privacy extensions for SLAAC to construct the autoconfig address, if the autoconfig is enabled.
Default	No IP address available, auto config is enabled	
Configuration Mode	config interface management	
History	3.1.0000	
Role	admin	
Example	`switch (config interface mgmt0) # ipv6 fe80::202:c9ff:fe5e:a5d8/64`	
Related Commands	ipv6 enable show interface <ifname>	
Notes	• On a given interface, up to 16 addresses can be configured • For Ethernet, the default interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface	

ipv6 dhcp primary-intf

ipv6 dhcp primary-intf <if-name>
no ipv6 dhcp primary-intf

Sets the interface from which non-interface-specific (resolver) configuration is accepted via DHCPv6.
The no form of the command resets non-interface-specific (resolver) configuration.

Syntax Description	if-name	Interface name: • lo • mgmt0 • mgmt1
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ipv6 dhcp primary-intf mgmt0` `switch (config) #`	
Related Commands	ipv6 enable ipv6 address show interface <ifname>	
Notes		

ipv6 dhcp stateless

	ipv6 dhcp stateless **no ipv6 dhcp stateless** Enables stateless DHCPv6 requests. The no form of the command disables stateless DHCPv6 requests.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # ipv6 dhcp stateless` `switch (config) #`
Related Commands	ipv6 enable ipv6 address show interface <ifname>
Notes	• This command only gets DNS configuration, not an IPv6 address • The no form of the command requests all information, including an IPv6 address

ipv6 dhcp client enable

	ipv6 dhcp client enable **[no] ipv6 dhcp client enable** Enables DHCPv6 on this interface. The no form of the command disables DHCPv6 on this interface.
Syntax Description	N/A
Default	ipv6 dhcp client enable
Configuration Mode	config interface management
History	3.7.11xx
Role	Admin
Example	`switch (config interface mgmt0) # ipv6 dhcp client enable`
Related Commands	ipv6 dhcp client renew show ipv6 dhcp
Notes	

ipv6 dhcp client renew

	ipv6 dhcp client renew Renews DHCPv6 lease for this interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface management
History	3.7.11xx
Role	Admin
Example	`switch (config interface mgmt0) # ipv6 dhcp client renew`
Related Commands	ipv6 dhcp client enable show ipv6 dhcp
Notes	

show interfaces mgmt0

	show interface mgmt0 Displays information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
	3.6.8008 Updated Example
Role	admin

Example

```
switch (config) # show interfaces mgmt0

Interface mgmt0 status:
  Comment         :
  Admin up        : yes
  Link up         : yes
  DHCP running    : yes
  IP address      : 10.12.67.33
  Netmask         : 255.255.255.128
  IPv6 enabled    : yes
  Autoconf enabled: no
  Autoconf route  : yes
  Autoconf privacy: no
  DHCPv6 running  : yes (but no valid lease)
  IPv6 addresses  : 1

  IPv6 address:
    fe80::268a:7ff:fe53:3d8e/64

  Speed           : 1000Mb/s (auto)
  Duplex          : full (auto)
  Interface type  : ethernet
  Interface source: bridge
  MTU             : 1500
  HW address      : 24:8A:07:53:3D:8E

  Rx:
    2055054  bytes
      28830  packets
          0  mcast packets
          0  discards
          0  errors
          0  overruns
          0  frame

  Tx:
    377716  bytes
      3200  packets
         0  discards
         0  errors
         0  overruns
         0  carrier
         0  collisions
         0  queue len
```

Related Commands N/A

Notes

show interfaces mgmt0 brief

	show interface mgmt0 brief Displays brief information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.8008 Updated Example
Role	admin

Example	`switch (config) # show interfaces mgmt0 brief`
	
	`Interface mgmt0 status:`
	`  Comment  :`
	`  Admin up  : yes`
	`  Link up  : yes`
	`  DHCP running  : yes`
	`  IP address  : 10.12.67.33`
	`  Netmask  : 255.255.255.128`
	`  IPv6 enabled  : yes`
	`  Autoconf enabled: no`
	`  Autoconf route  : yes`
	`  Autoconf privacy: no`
	`  DHCPv6 running  : yes (but no valid lease)`
	`  IPv6 addresses  : 1`
	
	`  IPv6 address:`
	`  fe80::268a:7ff:fe53:3d8e/64`
	
	`  Speed  : 1000Mb/s (auto)`
	`  Duplex  : full (auto)`
	`  Interface type  : ethernet`
	`  Interface source: bridge`
	`  MTU  : 1500`
	`  HW address  : 24:8A:07:53:3D:8E`

Related Commands	N/A
Notes	

show interfaces mgmt0 configured

show interface mgmt0 configured

Displays configuration information about the specified interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
	3.5.1000 Updated Example with “DHCP Hostname”
	3.6.8008 Updated Example
Role	admin

Example	`switch (config) # show interfaces mgmt0 configured`
	
	`Interface mgmt0 configuration:`
	`  Comment  :`
	`  Enabled  : yes`
	`  DHCP  : yes`
	`  DHCP Hostname  : yes`
	`  Zeroconf  : no`
	`  IP address  :`
	`  Netmask  :`
	`  IPv6 enabled  : yes`
	`  Autoconf enabled: no`
	`  Autoconf route  : yes`
	`  Autoconf privacy: no`
	`  DHCPv6 enabled  : yes`
	`  IPv6 addresses  : 0`
	`  Speed  : auto`
	`  Duplex  : auto`
	`  MTU  : 1500`

Related Commands	N/A
Notes	

4.1.7.2 Hostname Resolution

hostname

hostname <hostname>
no hostname

Sets a static system hostname.
The no form of the command clears the system hostname.

Syntax Description	hostname	A free-form string.
Default	Default hostname	
Configuration Mode	config	
History	3.1.0000	
	3.6.3004	Added support for the character “.”
Role	admin	
Example	`switch (config) # hostname my-switch-hostname`	
Related Commands	show hosts	
Notes	• Hostname may contain letters, numbers, periods (‘.’), and hyphens (‘-’), in any combination • Hostname may be 1-63 characters long • Hostname may not begin with a hyphen • Hostname may not contain other characters, such as “%”, “_” etc. • Hostname may not be set to one of the valid logging commands (i.e. debug-files, fields, files, format, level, local, monitor, receive, trap) • Changing the hostname stamps a new HTTPS certificate	

ip name-server

ip name-server <IPv4/IPv6 address>
no name-server <IPv4/IPv6 address>

Sets the static name server.
The no form of the command clears the name server.

Syntax Description	IPv4/v6 address	IPv4 or IPv6 address.
Default	No server name	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ip name-server 9.9.9.9`	
Related Commands	show hosts	
Notes		

ip domain-list

ip domain-list <domain-name>
no ip domain-list <domain-name>

Sets the static domain name.
The no form of the command clears the domain name.

Syntax Description	domain-name	The domain name in a string form. A domain name is an identification string that defines a realm of administrative autonomy, authority, or control in the Internet. Domain names are formed by the rules and procedures of the Domain Name System (DNS).
Default	No static domain name	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ip domain-list mydomain.com`	
Related Commands	show hosts	
Notes		

ip/ipv6 host

{ip | ipv6} host <hostname> <IP Address>
no {ip | ipv6} host <hostname> <IP Address>

Configures the static hostname IPv4 or IPv6 address mappings.
The no form of the command clears the static mapping.

Syntax Description	hostname	The hostname in a string form.
	IP Address	The IPv4 or IPv6 address.
Default	No static domain name.	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ip host my-host 2.2.2.2` `switch (config) # ipv6 host my-ipv6-host  2001::8f9`	
Related Commands	show hosts	
Notes		

ip/ipv6 map-hostname

	{ip |ipv6} map-hostname **no {ip | ipv6} map-hostname** Maps between the currently-configured hostname and the loopback address 127.0.0.1. The no form of the command clears the mapping.
Syntax Description	N/A
Default	IPv4 mapping is enabled by default IPv6 mapping is disabled by default
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # ip map-hostname`
Related Commands	show hosts
Notes	• If no mapping is configured, a mapping between the hostname and the IPv4 loop-back address 127.0.0.1 will be added • The no form of the command maps the hostname to the IPv6 loopback address if there is no statically configured mapping from the hostname to an IPv6 address (disabled by default) • Static host mappings are preferred over DNS results. As a result, with this option set, you will not be able to look up your hostname on your configured DNS server; but without it set, some problems may arise if your hostname cannot be looked up in DNS.

show hosts

	show hosts Displays hostname, DNS configuration, and static host mappings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show hosts`
	`Hostname: my-host-name`
	`Name server: 9.9.9.9 (configured)`
	`Name server: 11.11.11.11 (dynamic)`
	`Name server: 12.12.12.12 (dynamic)`
	`Name server: 13.13.13.13 (dynamic)`
	`Domain name: mydomain.com (configured)`
	`Domain name: example1.com (dynamic)`
	`Domain name: example2.com (dynamic)`
	`Domain name: example3.com (dynamic)`
	`Domain name: example4.com (dynamic)`
	`IP 1.1.1.1 maps to hostname p`
	`IP 3.3.3.3 maps to hostname localhost`
	`IP 2.2.2.2 maps to hostname my-host`
	`IPv6 ::1 maps to hostname localhost6`
	`Automatically map hostname to loopback address: yes`
	`Automatically map hostname to IPv6 loopback address: no`

Related Commands	N/A
Notes	

4.1.7.3 Routing

{ip | ipv6} route

{ip | ipv6} route [vrf <vrf-name>] {<network-prefix> <netmask> | <network -prefix>/<masklen>} <next-hop>
no ip route [vrf <vrf-name>] {<network-prefix> <netmask> | <network-prefix>/<masklen>} <next-hop>

Sets a static route for a given IP.
The no form of the command deletes the static route.

Syntax Description	network-prefix	IPv4 or IPv6 network prefix.
	netmask	IPv4 netmask formats are: • /24 • 255.255.255.0 IPv6 netmask format is: • /48 (as a part of the network prefix)
	nexthop-address	The IPv4 or IPv6 address of the next hop router for this route.
	ifname	The interface name (e.g., mgmt0, mgmt1).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ip route 20.20.20.0 255.255.255.0 mgmt0`	
Related Commands	show ip route	
Notes		

ipv6 default-gateway

ipv6 default-gateway {<ip-address> | <ifname>}
no ipv6 default-gateway

Sets a static default gateway.
The no form of the command deletes the default gateway.

Syntax Description	ip address	The default gateway IP address (IPv6).
	ifname	The interface name (e.g., mgmt0, mgmt1).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.2.0500	removed IPv4 configuration option
Role	admin	
Example	`switch (config) # ipv6 default-gateway ::1`	
Related Commands	show ip/ipv6 route show ipv6 default-gateway	
Notes	• The configured default gateway will not be used if DHCP is enabled. • In order to configure ipv4 default-gateway use ‘ip route’ command.	

show ip/ipv6 route

show {ip | ipv6} route [static]

Displays the routing table in the system.

Syntax Description	static	Filters the table with the static route entries.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show ip route`
	`Destination  Mask  Gateway  Interface  Source`
	`default  0.0.0.0  172.30.0.1  mgmt0  DHCP`
	`10.10.10.10  255.255.255.255  0.0.0.0  mgmt0  static`
	`20.10.10.10  255.255.255.255  172.30.0.1  mgmt0  static`
	`20.20.20.0  255.255.255.0  0.0.0.0  mgmt0  static`
	`172.30.0.0  255.255.0.0  0.0.0.0  mgmt0  interface`
	`switch (config) # show ipv6 route`
	`Destination prefix`
	`Gateway  Interface  Source`
	`----------`
	`::/0`
	`::  mgmt0  static`
	`::1/128`
	`::  lo  local`
	`2222:2222:2222::/64`
	`::  mgmt1  interface`	

Related Commands	ip route	
Notes		

show ipv6 default-gateway

show ipv6 default-gateway [static]

Displays the default gateway.

Syntax Description	static	Displays the static configuration of the default gateway
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # ipv6 default-gateway 10.10.10.10`
	`switch (config) # show ipv6 default-gateway`
	`Active default gateways:`
	`  172.30.0.1 (interface: mgmt0)`
	`switch (config) # show ipv6 default-gateway static`
	`Configured default gateway: 10.10.10.10`	

Related Commands	ipv6 default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enabled.	

4.1.7.4 Network to Media Resolution (ARP & NDP)

IPv4 network use Address Resolution Protocol (ARP) to resolve IP address to MAC address, while IPv6 network uses Network Discovery Protocol (NDP) that performs basically the same as ARP.

ip arp

ip arp <ip-address> <mac-address>
no ip arp <ip-address> <mac-address>

Sets a static ARP entry.
The no form of the command deletes the static ARP.

Syntax Description	IP address	IPv4 address.
	MAC address	MAC address.
Default	N/A	
Configuration Mode	config interface management	
History	3.2.0500	
Role	admin	
Example	`switch (config interface mgmt0) #ip arp 20.20.20.20 aa:aa:aa:aa:aa:aa`	
Related Commands	show ip arp ip route	
Notes		

ip arp timeout

ip arp [vrf <vrf-name>] timeout <timeout-value>
no ip arp [vrf <vrf-name>] timeout

Sets the dynamic ARP cache timeout.
The no form of the command sets the timeout to default.

Syntax Description	timeout-value	Time (in seconds) that an entry remains in the ARP cache. Range: 60-28800.
	vrf-name	VRF session name
Default	1500 seconds	
Configuration Mode	config	
History	3.2.0230	
	3.5.1000	Added VRF parameter and updated Notes
Role	admin	
Example	`switch (config) # ip arp timeout 2000` `switch (config) #`	
Related Commands	ip arp show ip arp	
Notes	• This value is used as the default ARP timeout whenever a new IP interface is created • The time interval after which each ARP entry becomes stale may actually vary from 50-150% of the configured value	

show ip arp

show ip arp [interface <type> | <ip-address> | count]

Displays ARP table.

Syntax Description	interface type	Filters the table according to a specific interface (i.e. mgmt0)
	ip-address	Filters the table to the specific ip-address
	count	Shows ARP statistics
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3000	
Role	admin	

Example	`switch-626a54 [standalone: master] (config) # show ip arp`
	
	`Total number of entries: 3`
	
	`  Address             Type           Hardware Address         Interface`
	`  ------------------------------------------------------------------------`
	`---`
	`  10.209.0.1           Dynamic ETH     00:00:5E:00:01:01         mgmt0`
	`  10.209.1.120         Dynamic ETH     00:02:C9:62:E8:C2         mgmt0`
	`  10.209.1.121         Dynamic ETH     00:02:C9:62:E7:42         mgmt0`
	`switch (config) # show ip arp count`
	`ARP Table size: 3 (inband: 0, out of band: 3)`
	`switch (config) #`	

Related Commands		
Notes		

ipv6 neighbor

ipv6 neighbor <ipv6-address> <ifname> <mac-address>
no ipv6 neighbor <ipv6-address> <ifname> <mac-address>

Adds a static neighbor entry.
The no form of the command deletes the static entry.

Syntax Description	IPv6 address	The IPv6 address
	ifname	The management interface (i.e. mgmt0, mgmt1)
	MAC address	The MAC address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ipv6 neighbor 2001:db8:701f::8f9 mgmt0 00:11:22:33:44:55` `switch (config) #`	
Related Commands	show ipv6 neighbor ipv6 route arp clear ipv6 neighbors	
Notes	• ARP is used only with IPv4. In IPv6 networks, Neighbor Discovery Protocol (NDP) is used similarly. • Use The no form of the command to remove static entries. Dynamic entries can be cleared via the “clear ipv6 neighbors” command.	

clear ipv6 neighbors

clear ipv6 neighbors {ethernet <port> | vlan <vlan-id> | port-channel <id> | vrf <vrf-id>} [<ip-addr>]

Clears the dynamic neighbors cache.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.4110 Updated command.
Role	admin
Example	`switch (config) # clear ipv6 neighbors` `switch (config) #`
Related Commands	ipv6 neighbor show ipv6 neighbor arp
Notes	• Clearing Neighbor Discovery Protocol (NDP) cache removes only the dynamic entries learned and not the static entries configured • Use the no form of the command to remove static entries

show ipv6 neighbors

show ipv6 neighbors [static]

Displays the Neighbor Discovery Protocol (NDP) table.

Syntax Description	static	Filters only the table of the static entries.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show ipv6 neighbors`
	`IPv6 Address Age MAC Address State Interf`
	`---`
	`2001::2 9428 AA:AA:AA:AA:AA:AA permanent mgmt0`
	`switch (config) #`	

Related Commands	ipv6 neighbor clear ipv6 neighbor show ipv6	
Notes		

4.1.7.5 DHCP

ip dhcp

ip dhcp {default-gateway yield-to-static| hostname <hostname>| primary-intf <ifname> | send-hostname }
no ip dhcp {default-gateway yield-to-static| hostname | | primary-intf | send-hostname}

Sets global DHCP configuration.
The no form of the command deletes the DHCP configuration.

Syntax Description	yield-to-static|	Does not allow you to install a default gateway from DHCP if there is already a statically configured one.
	hostname	Specifies the hostname to be sent during DHCP client negotiation if send-hostname is enabled.
	primary-intf <ifname>	Sets the interface from which a non-interface-specific configuration (resolver and routes) will be accepted via DHCP.
	send-hostname	Enables the DHCP client to send a hostname during negotiation.
Default	no ip dhcp yield-to-static no ip dhcp hostname ip ip dhcp primary-intf mgmt0 no ip dhcp send-hostname	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # ip dhcp default-gateway yield-to-static`	
Related Commands	show ip dhcp dhcp [renew]	
Notes	DHCP is supported for IPv4 networks only.	

show ip dhcp

	show ip dhcp Displays the DHCP configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.5000 Updated Example
Role	admin

Example	`switch (config) # show ip dhcp`
	`----------`
	`Interface DHCP DHCP Valid`
	`Enabled Running lease`
	`----------`
	`dummy0 no no no`
	`lo no no no`
	`mgmt0 yes yes yes`
	`mgmt1 no no no`
	`mgmts0 no no no`
	`mgmts1 no no no`
	`vif1 no no no`
	
	`IPv4 dhcp default gateway yields to static configuration: no`
	
	`DHCP primary interface:`
	`Configured: mgmt0`
	`Active: mgmt0`
	
	`DHCP client options:`
	`Send Hostname: no`
	`Client Hostname: 1.1.1.1`

Related Commands	ip dhcp dhcp [renew]
Notes	

4.1.7.6 General IPv6 Commands

ipv6 enable

	ipv6 enable **no ipv6 enable** Enables IPv6 globally on the management interface. The no form of the command disables IPv6 globally on the management interface.
Syntax Description	N/A
Default	IPv6 is disabled
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # ipv6 enable`
	`switch (config) # show ipv6`
	`IPv6 summary`
	`    IPv6 supported:        yes`
	`    IPv6 admin enabled:    yes`
	`    IPv6 interface count:  2`
	`switch (config) #`

Related Commands	ipv6 default-gateway ipv6 host ipv6 map-hostname ipv6 neighbor ipv6 route show ipv6 show ipv6 default-gateway show ipv6 route
Notes	

4.1.7.7 IP Diagnostic Tools

ping

ping [-LRUbdfnqrvVaA] [-c count] [-i interval] [-w deadline] [-p pattern] [-s packetsize] [-t ttl] [-I interface or address] [-M mtu discovery hint] [-S sndbuf] [-T timestamp option] [-Q tos] [hop1 ...] destination

Sends ICMP echo requests to a specified host.

Syntax Description	Linux Ping options	https://www.lifewire.com/uses-of-command-ping-2201076
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # ping 172.30.2.2`
	`PING 172.30.2.2 (172.30.2.2) 56(84) bytes of data.`
	`64 bytes from 172.30.2.2: icmp_seq=1 ttl=64 time=0.703 ms`
	`64 bytes from 172.30.2.2: icmp_seq=2 ttl=64 time=0.187 ms`
	`64 bytes from 172.30.2.2: icmp_seq=3 ttl=64 time=0.166 ms`
	`64 bytes from 172.30.2.2: icmp_seq=4 ttl=64 time=0.161 ms`
	`64 bytes from 172.30.2.2: icmp_seq=5 ttl=64 time=0.153 ms`
	`64 bytes from 172.30.2.2: icmp_seq=6 ttl=64 time=0.144 ms`
	`^C`
	`--- 172.30.2.2 ping statistics ---`
	`6 packets transmitted, 6 received, 0% packet loss, time 5004ms`
	`rtt min/avg/max/mdev = 0.144/0.252/0.703/0.202 ms`
	`switch (config) #`	

Related Commands	traceroutes	
Notes		

traceroute

traceroute [-46dFITUnrAV] [-f first_ttl] [-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] host [packetlen]

Traces the route packets take to a destination.

Syntax Description		
	-4	Uses IPv4
	-6	Uses IPv6
	-d	Enables socket level debugging
	-F	Sets DF (do not fragment bit) on
	-I	Uses ICMP ECHO for tracerouting
	-T	Uses TCP SYN for tracerouting
	-U	Uses UDP datagram (default) for tracerouting
	-n	Does not resolve IP addresses to their domain names
	-r	Bypasses the normal routing and send directly to a host on an attached network
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses
	-V	Prints version info and exit
	-f	Starts from the first_ttl hop (instead from 1)
	-g	Routes packets through the specified gateway (maximum 8 for IPv4 and 127 for IPv6)
	-i	Specifies a network interface with which to operate
	-m	Sets the max number of hops (max TTL to be reached). Default is 30
	-N	Sets the number of probes to be tried simultaneously (default is 16)
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets
	-l	Uses specified flow_label for IPv6 packets

-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
-q	Sets the number of probes per each hop. Default is 3.
-s	Uses source src_addr for outgoing packets.
-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).

Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # traceroute 192.168.10.70`
	`traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte packets`
	`1 172.30.0.1 (172.30.0.1) 3.632 ms 2.849 ms 3.544 ms`
	`2 10.222.128.46 (10.222.128.46) 3.176 ms 3.289 ms 3.656 ms`
	`3 10.158.128.30 (10.158.128.30) 15.331 ms 15.819 ms 16.388 ms`
	`4 10.158.128.65 (10.158.128.65) 20.468 ms 7.893 ms 12.27 ms`
	`5 10.7.34.115 (10.7.34.115) 16.405 ms 11.985 ms 12.264 ms`
	`6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms`
	`switch (config) #`

Related Commands	
Notes	

tcpdump

tcpdump [-aAdDeflLnNOpqRStuUvxX] [-c count] [-C file_size]
[-E algo:secret] [-F file] [-i interface] [-M secret]
[-r file] [-s snaplen] [-T type] [-w file]
[-W filecount] [-y datalinktype] [-Z user]
[-D list possible interfaces] [expression]

Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # tcpdump`
	`......`
	`09:37:38.678812 IP 192.168.10.7.ssh > 192.168.10.1.54155: P`
	`1494624:1494800(176) ack 625 win 90`
	`<nop,nop,timestamp 5842763 858672398>`
	`09:37:38.678860 IP 192.168.10.7.ssh > 192.168.10.1.54155: P`
	`1494800:1495104(304) ack 625 win 90`
	`<nop,nop,timestamp 5842763 858672398>`
	`...`
	`9141 packets captured`
	`9142 packets received by filter`
	`0 packets dropped by kernel`
	`switch (config) #`

Related Commands	N/A
Notes	

4.2 Management Source IP Address

In many cases network operators prefer to have a single IP address for the switch that is used for management operations like switch configuration, receiving remote log files, ping, etc. The single IP address is needed for building firewall rules, so network switches can be easily identified. It is also required for identifying management traffic and exact management target in network logs.

The following protocols are supported by the feature:

- FTP
- TFTP
- NTP
- Syslog
- TACACS
- SSH, SSHD, SCP
- Ping
- Traceroute

4.2.1 Commands

ssh server listen

ssh server listen <interface>
[no] ssh server listen <interface>

Defines a source interface for ssh server.

Syntax Description	interface	Interface to bind. May be mgmt. 0, lo or loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	
Example	N/A	
Related Commands	N/A	
Notes		

ssh client global source-interface

ssh client global source-interface <interface>
[no] ssh client global source-interface <interface>

Binds SSH client to specific address used by the slogin command.

Syntax Description	interface	Interface to bind. May be loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	
Example	N/A	
Related Commands	N/A	
Notes		

ip ftp source-interface

ip ftp source-interface <interface>
[no] ip ftp source-interface <interface>

Configures the source interface.

Syntax Description	interface	Interface to bind. May be loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	

Example	`switch (config)# no ip ftp source-interface`
	`switch (config)# ip ftp source-interface loopback7`
	`switch (config)# show ip ftp source-interface`
	
	`Source IP for ftp client:`
	`  Configured: loopback7`
	`  Current   : loopback7`
	`  IPv4-addr : 5.5.5.5`
	`  IPv6-addr : none`	

Related Commands	N/A	
Notes		

ip tftp source-interface

ip tftp source-interface <interface>
[no] ip tftp source-interface <interface>

Configures the source interface.

Syntax Description	interface	Interface to bind. May be loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	

Example	`switch (config)# no ip tftp source-interface`
	`switch (config)# ip tftp source-interface loopback7`
	`switch (config)# show ip tftp source-interface`
	
	`Source IP for tftp client:`
	`  Configured: loopback7`
	`  Current   : loopback7`
	`  IPv4-addr : 5.5.5.5`
	`  IPv6-addr : none`	

Related Commands	N/A	
Notes		

logging source-interface

logging source-interface <interface>
[no] logging source-interface <interface>

Configures the source interface.

Syntax Description	interface	Interface to bind. May be loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	

Example	`switch (config)# no logging source-interface`
	`switch (config)# logging source-interface loopback7`
	`switch (config)# show logging source-interface`
	
	`Source IP for syslogd client:`
	`  Configured: loopback7`
	`  Current   : loopback7`
	`  IPv4-addr : 5.5.5.5`
	`  IPv6-addr : none`	

Related Commands	N/A	
Notes		

tacacs source-interface

tacacs source-interface <interface>
[no] tacacs source-interface <interface>

Configures the source interface.

Syntax Description	interface	Interface to bind. May be loopback 0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	
Example	N/A	
Related Commands	N/A	
Notes		

ip icmp source-interface

	ip icmp source-interface **[no] ip icmp source-interface** Configures the source interface for ping command.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin
Example	N/A
Related Commands	N/A
Notes	

ntp source-interface

ntp source-interface <interface>
[no] ntp source-interface <interface>

Configures the source interface.

Syntax Description	interface	Interface to bind. May be loopback0-31.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	

Example	`switch (config)# no ntp source-interface`
	`switch (config)# ntp source-interface loopback7`
	`switch (config)# show ntp source-interface`
	
	`Source IP for ntp client:`
	`  Configured: loopback7`
	`  Current   : loopback7`
	`  IPv4-addr : 5.5.5.5`
	`  IPv6-addr : none`	

Related Commands	N/A	
Notes	This command sets source ip for ntpd and ntpdate.	

show ip ftp source-interface

show ip ftp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for ftp client:`
	`  Configured: loopback7`
	`  Current   : loopback7`
	`  IPv4-addr : 5.5.5.5`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	

show ntp source-interface

show ntp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for ntp client:`
	`  Configured: loopback2`
	`  Current   : loopback2`
	`  IPv4-addr : 10.7.144.97`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show logging source-interface

	show logging source-interface Shows the source interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for syslogd client:`
	`  Configured: loopback23`
	`  Current   : loopback23`
	`  IPv4-addr : 1.3.5.7`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	

show tacacs source-interface

show tacacs source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for tacacs client:`
	`  Configured: none`
	`  Current   : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	

show ntp source-interface

show ntp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for ntp client:`
	`  Configured: loopback2`
	`  Current   : loopback2`
	`  IPv4-addr : 10.7.144.97`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show ip icmp source-interface

show ip icmp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for ping client:`
	`  Configured: none`
	`  Current   : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show ip traceroute source-interface

show ip traceroute source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for traceroute client:`
	`  Configured: none`
	`  Current   : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show ssh client source-interface

	show ssh client source-interface Shows the ssh client source interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000 3.7.11xx Updated command syntax & example
Role	Admin

Example	`(config) # show ssh client source-interface`
	`Source IP for ssh client:`
	`  Configured: loopback1`
	`  Current   : loopback1`
	`  IPv4-addr : 1.1.1.1`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show ip scp source-interface

show ip scp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for scp client:`
	`  Configured: none`
	`  Current   : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

show ip sftp source-interface

show ip sftp source-interface

Shows the source interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	admin

Example	`Source IP for sftp client:`
	`  Configured:`
	`  Current   : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`

Related Commands	N/A
Notes	N/A

4.3 NTP, Clock & Time Zones

Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks. NTP is intended to synchronize all participating computers to within a few milliseconds of Coordinated Universal Time (UTC) and is designed to mitigate the effects of variable network latency. NTP can usually maintain time to within tens of milliseconds over the public Internet, and can achieve better than one millisecond accuracy in local area networks under ideal conditions.

For an example, please refer to "HowTo enable NTP on Mellanox switches" in the Mellanox Community (https://community.mellanox.com).

4.3.1 NTP Authenticate

When authentication of incoming NTP packets is enabled, the switch ensures that they come from an authenticated time source before using them for time synchronization on the switch. Authentication keys are created and added to the trusted list.

➢ ***To add a key to be used for authentication***

Step 1. Create the key. Run:

```
switch (config)# ntp authentication-key 1 md5 password
```

Step 2. Add the key to the trusted list. Run:

```
switch (config)# ntp trusted-key 1
```

Step 3. Assign the key to the server/peer. Run:

```
switch (config)# ntp server 10.34.1.1 keyID 1
```

4.3.2 NTP Authentication Key

An authentication key may be created and used to authenticate incoming NTP packets.

For the key to be used:

1. It should be shared with the NTP server/peer sending the NTP packet.
2. It should be added to the trusted list.
3. NTP authenticate should be enabled on the switch.

4.3.3 Commands

clock set

clock set <hh:mm:ss> [<yyyy/mm/dd>]

Sets the time and date.

Syntax Description	hh:mm:ss	Time.
	yyyy/mm/dd	Date.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # clock set 23:23:23 2010/08/19`	
Related Commands	show clock	
Notes	If not specified, the date will be left the same.	

clock timezone

clock timezone [<zone word> [<zone word> [<zone word>] [<zone word>]]]

Sets the system time zone. The time zone may be specified in one of three ways:

- A nearby city whose time zone rules to follow. The system has a large list of cities which can be displayed by the help and completion system. They are organized hierarchically because there are too many of them to display in a flat list. A given city may be required to be specified in two, three, or four words, depending on the city.
- An offset from UTC. This will be in the form UTC-offset UTC, UTC-offset UTC+<0-14>, UTC-offset UTC-<1-12>.
- UTC (Universal Time, which is almost identical to GMT), and this is the default time zone

The no form of the command resets time zone to its default (GMT).

Syntax Description	zone word	The possible forms this could take include: continent, city, continent, country, city, continent, region, country, city, ocean, and/or island.
Default	GMT	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # clock timezone America North United_States Other New_York`	
Related Commands	show clock	
Notes		

ntp

ntp {disable | enable | {peer | server} <IP address> [version <number> | disable]}
no ntp {disable | enable | {peer | server} <IP address> [version <number> | disable]}

Configures NTP.
The no form of the command negates NTP options.

Syntax Description	disable	Disables NTP
	enable	Enables NTP
	peer or server	Configures an NTP peer or server node
	IP address	IPv4 or IPv6 address
	version <number>	Specifies the NTP version number of this peer Possible values: 3 or 4
Default	NTP is enabled NTP version number is 4	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # no ntp peer 192.168.10.24 disable` `switch (config) #`	
Related Commands	N/A	
Notes		

ntpdate

ntpdate <IP address>

Sets the system clock using the specified SNTP server.

Syntax Description	IP address IP.
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # ntpdate 192.168.10.10`
	`26 Feb 17:25:40 ntpdate[15206]: adjust time server 192.168.10.10 offset -0.000092 sec`
	`switch (config) #`

Related Commands	N/A
Notes	This is a one-time operation and does not cause the clock to be kept in sync on an ongoing basis. It will generate an error if SNTP is enabled since the socket it requires will already be in use.

ntp authenticate

	ntp authenticate **no ntp authenticate** Enables NTP authentication. The no form of the command disables NTP authentication.
Syntax Description	N/A N/A
Default	Disabled
Configuration Mode	config
History	3.5.0200
Role	admin
Example	`switch (config) # ntp authenticate`
Related Commands	N/A
Notes	

ntp authentication-key

ntp authentication-key <key_id> <encrypt_type> [<password>]
no ntp authentication-key <key_id>

Adds a new authentication key and stores it.
The no form of the command removes key ID configuration if it exists.

Syntax Description	key_id	Specifies a key ID, whether existing or a new one to be added. Range: 1-65534.
	encrypt_type	Specifies encryption type to use (md5, or sha1)
	password	Password string
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
Role	admin	

Example	`switch (config) # ntp authentication-key 123 md5 examplepass`
	`switch (config) # ntp authentication-key 1234 sha1`
	`Password: **`
	`Confirm: **`
	`switch (config) #`	

Related Commands	N/A	
Notes	If a password is not entered, a prompt appears requiring that a password is introduced.	

ntp peer disable

ntp peer <ip_address> disable
no ntp peer <ip_address> disable

Temporarily disables this NTP peer.
The no form of the command enables this NTP peer.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
	3.6.4000	Added hostname as option for ip_address, and added Notes.
Role	admin	
Example	`switch (config) # ntp peer 10.10.10.10 disable` `switch (config) #`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp peer keyID

ntp peer <ip_address> keyID <key_id>
no ntp peer <ip_address> keyID <key_id>

Specifies the KeyID of the NTP peer.
The no form of the command removes key ID configuration from the NTP peer.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
	key_id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp peer 10.10.10.10 keyID 120`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp peer version

ntp peer <ip_address> version <ver_num>
no ntp peer <ip_address> version <ver_num>

Specifies the NTP version number of this peer.
The no form of the command defaults NTP to version 4.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
	ver_num	NTP version (3 or 4)
Default	4	
Configuration Mode	config	
History	3.5.0200	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp peer 10.10.10.10 version 4`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp server disable

ntp server <ip_address> disable
no ntp server <ip_address> disable

Temporarily disables this NTP server.
The no form of the command enables this NTP server.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	Disabled	
Configuration Mode	config	
History	3.5.0000	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp server 10.10.10.10 disable` `switch (config) #`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp server keyID

ntp server <ip_address> keyID <key_id>
no ntp server <ip_address> keyID <key_id>

Specifies the KeyID of the NTP server.
The no form of the command removes key ID configuration from the NTP server.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
	key_id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp server 10.10.10.10 keyID 120`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp server trusted-enable

ntp server <ip_address> trusted-enable
no ntp server <ip_address> trusted-enable

Trusts this NTP server; if authentication is configured this will additionally force all time updates to only use trusted servers.
The no form of the command removes trust from this NTP server

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	N/A	
Configuration Mode	config	
History	3.6.2002	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp server 10.10.10.10 trusted-enable`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. • NTP trusted servers can be used as a mitigation for Sybil attacks which is a vulnerability caused by NTP peers sharing the same NTP key base. This mitigation adds the concept of trusted servers which if enabled in conjunction with NTP authentication ensures that time information will only be obtained from trusted servers.	

ntp server version

ntp server <ip_address> version <ver_num>
no ntp server <ip_address> version <ver_num>

Specifies the NTP version number of this server.
The no form of the command defaults NTP to version 4.

Syntax Description	ip_address	IP address of the peer. IPv4, IPv6 and hostname (FQDN) are acceptable.
	ver_num	NTP version (3 or 4)
Default	4	
Configuration Mode	config	
History	3.5.0200	
	3.6.4000	Added hostname as ip_address option and added Notes.
Role	admin	
Example	`switch (config) # ntp server 10.10.10.10 version 4`	
Related Commands	N/A	
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone id for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

ntp trusted-key

ntp trusted-key <key(s)>
no ntp trusted-key <key(s)>

Adds one or more keys to the trusted key list.
The no form of the command removes keys from the trusted key list.

Syntax Description	key(s)	Range: 1-65534.
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # ntp trusted-key 1,3,5` `switch (config) # ntp trusted-key 1-5`	
Related Commands	ntp authentication-key	
Notes	Keys may be separated with commas without any space, or they may be set as a range using a hyphen.	

show clock

	show clock Displays the current system time, date and time zone.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show clock`
	`Time:  02:48:41`
	`Date:  2018/1/1`
	`Time zone:  UTC (Etc/UTC)`
	`UTC offset:  same as UTC`

Related Commands	N/A
Notes	

show ntp

	show ntp Displays the current NTP settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.0200 Updated Example 3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ntp`
	`NTP is administratively enabled.`
	`NTP Authentication is administratively disabled.`
	`Clock is synchronized.  Reference: 108.61.73.244.  Offset: -2.833 ms.`
	`Active servers and peers:`
	
	`108.61.73.244  # Hostname configuration`
	`  Configured as  : 0.us.pool.ntp.org`
	`  Conf Type  : server`
	`  Status  : sys.peer(*)`
	`  Stratum  : 2`
	`  Offset(msec)  : -2.833`
	`  Ref clock  : 128.59.0.245`
	`  Poll Interval (sec): 256`
	`  Last Response (sec): 203`
	`  Auth state  : none`
	
	`10.7.144.19  # IP configuration`
	`  Conf Type  : peer`
	`  Status  : sys.peer(*)`
	`  Stratum  : 2`
	`  Offset(msec)  : -1.747`
	`  Ref clock  : 128.59.0.245`
	`  Poll Interval (sec): 64`
	`  Last Response (sec): 1`
	`  Auth state  : none`

Related Commands	N/A
Notes	

show ntp configured

	show ntp configured Displays NTP configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200 3.6.6102 Updated Example
Role	admin

Example	`switch (config) # show ntp configured`
	`NTP enabled: yes`
	`NTP Authentication enabled: no`
	`NTP peer 0.us.pool.ntp.org  # Hostname peer configuration`
	`  Resolved as: 45.79.111.114`
	`  Enabled: yes`
	`  NTP version: 4`
	`  Key ID: none`
	`NTP peer 2.3.1.3  # IP peer configuration`
	`  Enabled: yes`
	`  NTP version: 4`
	`  Key ID: none`
	`NTP server vnc23  # Hostname server configuration`
	`  Resolved as: 10.7.2.23`
	`  Enabled: yes`
	`  NTP version: 4`
	`  Key ID: none`
	`  Trusted: no`
	`NTP server 1.2.3.4  # IP server configuration`
	`  Enabled: yes`
	`  NTP version: 4`
	`  Key ID: none`
	`  Trusted: no`
	`NTP server idontexist (DNS resolution failed. Reset or reconfigure NTP to try again)`
	`  Enabled: yes`
	`  NTP version: 4`
	`  Key ID: none`
	`  Trusted: no`

Related Commands	N/A
Notes	

show ntp keys

	show ntp configured Displays NTP keys.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200
Role	admin

Example	`switch (config) # show ntp keys`
	`NTP Key 1`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 2`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 3`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 4`
	`  Trusted: yes`
	`  Encryption Type: md5`
	`switch (config) #`

Related Commands	N/A
Notes	

4.4 Precision Time Protocol

Synchronizing network applications require their wall clock time to be aligned precisely with a reference time source (to the order of micro seconds or less). To achieve such accuracy, the application needs the support of networking HW (switch and adapter card), to provide the means to stamp time-sensitive packets. It also requires a time synchronization protocol which would make use of the HW time stamping to adjust its wall clock time to an accurate clock in the network.

4.4.1 PTP Principles

The basic principal of PTP is as follows: Slave time = master time + propagation delay + offset.

The purpose of the protocol is to align the slave and the master time so that the gap between them is the propagation delay of the packet. Or in other words, the purpose of the protocol is to use the offset to correct the slave time so the offset between the master sending the packet and the slave receiving the packet is the propagation delay.

Master time is sent periodically by a reliable clock source named Master Clock (MC). In a PTP network, one single reference source is elected called Grand Master Clock (GMC). Propagation delay is calculated between each node and the MC by one of the two methods provided by the standard and further explained below.

To reach sub-micro second resolutions, all the time stamps which record when a packet is sent and received should be done in the HW. This may impose interaction between SW and HW to query the HW time and send follow-up messages. This issue is further explained below in 2 step section.

Assuming that the propagation delay in the network is symmetric, the propagation time is the average time that took the sync and delay req messages to be switched.

Propagation delay = (T4-T1-(T3-T2))/2=(T4-T1+T2-T3)/2

T1 represents the time that the packet left the master which is actually the master time.

Figure 10 provides an example of the stages required by a slave clock to align its time to the master clock:

Figure 10: PTP Clock Synchronization Example

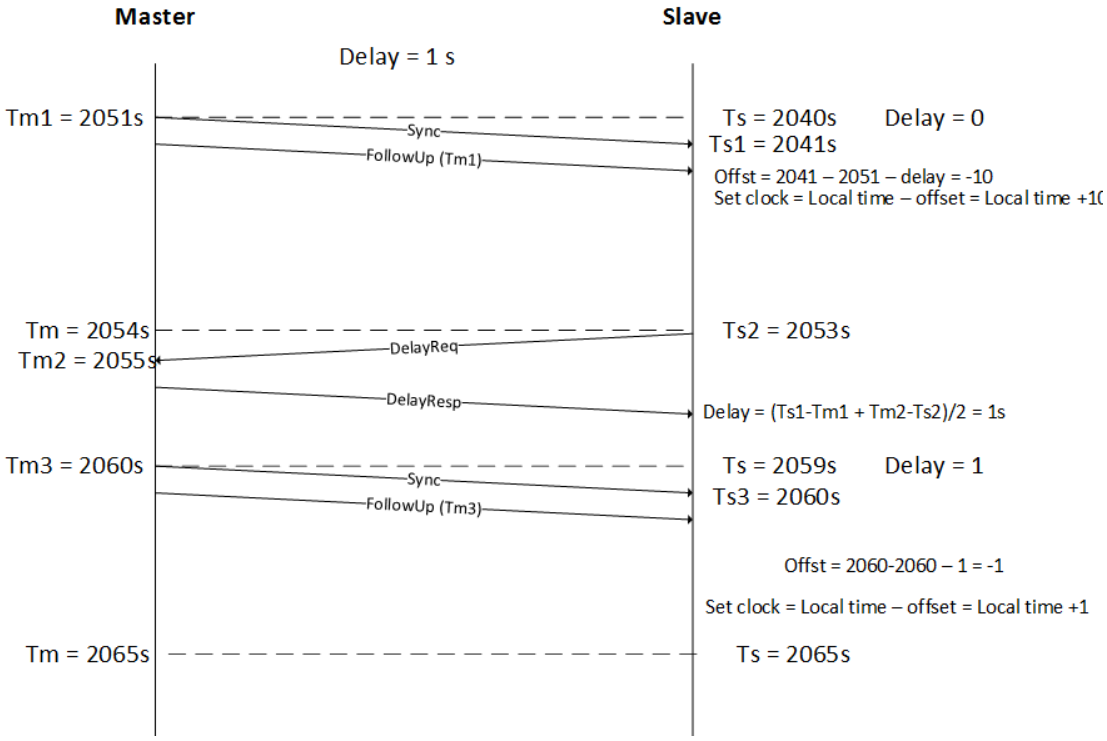

Table 24 - PTP Message Formats

Message Type	Hex Value	Class
Sync	0	Event
Follow-up	8	General
Delay_Req	1	Event
Delay_Resp	9	General
Pdelay_Req	2	Event
Pdelay_Resp	3	Event
Pdelay_Resp followup	A	General
Announce	B	General
Signaling	C	General
Management	D	General

4.4.2 Clock Types and Operation Modes

The types of clocks available are as follows:

- Grand Master Clock (GMC) – the reference time source derived from an accurate clock such as a GNSS driven clock (i.e. GPS, GLONASS, GALILEO)
- Boundary Clock (BC) – a network device that acts as slave to its master and as master to its slaves. (Mellanox implements only this)
- Ordinary Clock (OC) – a clock that operates either as a Master or a Slave. In the case of a slave, the end point whose clock is been synced (normally a host/server).
- Master Clock (MC) – a clock which operates as a Master and derives its timing capabilities from the clock chain up to the GMC. It typically serves as a port on a BC connected to a host running as a slave.
- Transparent Clock (TC) – a PTP aware switch capable of measuring the PTP packet switching delay (transient time) and updating the data in the packet. In peer-to-peer (P2P) delay calculation mechanism, a TC device is also required to calculate its delay from the next hop toward the MC and add the value to the switching delay.

Two modes of delay calculations are defined:

- End-to-End (E2E) – each slave calculates its delay from the MC by running Delay request/ delay response sequence (Mellanox implements only this)
- Peer-to-Peer – propagation delay (Pdelay) is calculated periodically on each link between the slave and the MC independently. The time synchronization packet sent from the MC to all the slaves in the network is updated by each of the downstream nodes with both switching delay (the time that the packet traversed the switch) and upstream hop Pdelay.

4.4.3 PTP Domains

A domain consists of one or more PTP devices communicating with each other. PTP domain defines the scope of PTP message communication, state, operations, data sets, and timescale.

4.4.3.1 Boundary Clock

In a full E2E PTP deployment, the GMC needs to respond to each slave's delay request message. A normal profile of PTP may require a few delay calculations per second. An average GMC is capable of addressing few thousands of messages per second. This imposes that direct slave/ GMC communication limits the number of overall OCs to ~8K. To scale beyond that, there is a need for a hierarchy between the GMC and the slave. This is achieved by implementing BC, either in the TOR switches or on all the switches in the DC.

Figure 11 shows the master/slave role that a boundary clock implements between the MC and the Slave (OC).

Figure 11: Boundary Clock Master/Slave Functionality

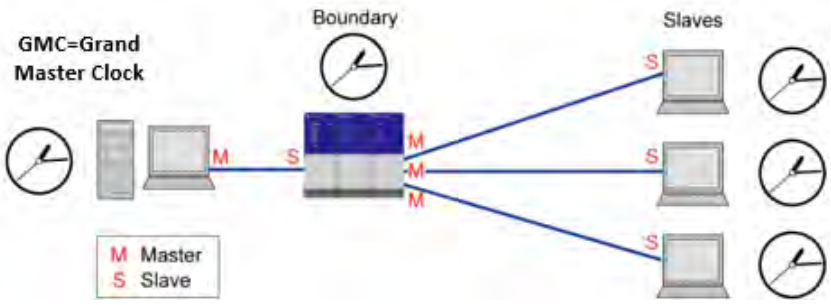

Each BC acts as a slave towards the GMC and as GMC to its local slaves. Although adding a BC device introduces accuracy degradation as explained above, it becomes mandatory when the number of slaves on a single MC exceeds few thousand devices.

Another use of BC is to bridge between networks. When running PTP over native Ethernet packets, to create larger PTP domains, there is a need to bridge between the broadcast domains. This is done by BC switches.

Table 25 - Default PTP Profile Attributes (SMPTE 2059-2)

Name	Range	Default
Announce interval	-3 (0.125s), 1 (2s)	-2 (0.25s)
Announce timeout interval	2, 10	3
Sync interval (logSyncInt)	-7, -1	-3
Delay request interval	logSyncInt, logSyncInt +5	logSyncInt
PTP domain	0, 127	127
Priority 1	0, 255	128
Priority 2	0, 255	128

4.4.4 Configuring PTP

IEEE 1588 Precision Time Protocol (PTP) may be configured either on router or switch interfaces.

To enable PTP on a router interface you could simply enable it on the selected interface.

The process of configuring PTP on a switch interface is slightly different, however. PTP should be enabled on the interface itself as well as on the respective VLAN interface(s).

All PTP configuration for switch interfaces is taken from those defined on the VLAN interface.

Prior to enabling PTP, NTP must be disabled.

➢ ***To configure PTP on a router interface:***

Step 1. Enable the PTP CLI commands. Run:

```
switch (config) # protocol ptp
```

Step 2. Configure the router interface. Run:

```
switch (config) #  1/1 no switchport force
```

Step 3. Add the primary IP address. Run:

```
switch (config) #  1/1 ip address 172.16.1.1/24
```

Step 4. Enable PTP on the interface. Run:

```
switch (config) #  1/1 ptp enable
```

➢ ***To verify the PTP configuration:***

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Enabled
Domain                  : 127
Clock identity          : 7C:FE:90:FF:FE:FA:21:88
GMC identity            : 7C:FE:90:FF:FE:FA:21:88
Number of master ports  : 1
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------
  Port           VLAN      State        Forced Master
  ------------------------------------------------
  Eth1/1         N/A       MASTER       no
```

➢ ***To configure PTP on a switch interface:***

Step 1. Enable the PTP CLI commands. Run:

```
switch (config) # protocol ptp
```

Step 2. Add the VLANs. Run:

```
switch (config) # vlan 2-3
```

Step 3. Configure VLAN membership.

For access interfaces, run:

```
switch (config) #  1/2 switchport mode access
switch (config) #  1/2 switchport access vlan 2
```

For trunked interfaces, run:

```
switch (config) #  1/1 switchport mode trunk
```

Step 4. Enable PTP on the VLAN interface. Run:

```
switch (config) # interface vlan 2 ptp enable
switch (config) # interface vlan 3 ptp enable
```

Step 5. Enable PTP on the interface. Run:

```
switch (config) #  1/1 ptp enable
```

The interface must be a member of the PTP enabled VLAN(s).

➢ ***To verify the PTP configuration:***

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Enabled
Domain                  : 127
Clock identity          : 7C:FE:90:FF:FE:FA:21:88
GMC identity            : 7C:FE:90:FF:FE:FA:21:88
Number of master ports  : 2
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------
  Port          VLAN      State        Forced Master
  ------------------------------------------------
  Eth1/1        2         MASTER       no
  Eth1/2        2         MASTER       no
  Eth1/1        3         SLAVE        no
```

4.4.5 Securing PTP Infrastructure

To protect the switch from rogue or mis-configured PTP endpoints, you may secure your Boundary Clock ports by creating an Acceptable Master Table (AMT) and configuring known PTP ports to always behave as a master port via the Forced Master option.

The AMT is a whitelist of up to 8 clock identities that are admissible to take part as valid Grand-Masters in the Best Master Clock Algorithm (BMCA).

The Forced Master is enabled on a per-port basis to prevent processing announce messages from a PTP endpoint connected to it, in order for it to always stay in a Master state.

To configure Forced Master on a switch interface, you must enable it on the interface itself as well as on the respective VLAN interface(s).

➢ ***To configure Acceptable Master Table:***

Add the validated clock identities:

```
switch (config) # ptp amt E4:1D:2D:FF:FE:46:13:88
switch (config) # ptp amt E4:1D:2D:FF:FE:44:23:B7
```

➢ ***To verify the Acceptable Master Table configuration:***

```
switch (config) # show ptp amt

Clock Identities:
  E4:1D:2D:FF:FE:44:23:B7
  E4:1D:2D:FF:FE:46:13:88
```

➢ ***To configure Forced Master on a router interface:***

To enable Forced Master on the interface:

```
switch (config) #  1/2 ptp enable forced-master
```

➢ ***To verify PTP configuration:***

```
switch (config) # show ptp
PTP mode               : Boundary Clock
Message format         : Mixed
Acceptable Master Table : Enabled
Domain                 : 127
Clock identity         : 7C:FE:90:FF:FE:FA:21:88
GMC identity           : 7C:FE:90:FF:FE:FA:21:88
Number of master ports  : 1
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------
  Port          VLAN      State        Forced Master
  ------------------------------------------------
  Eth1/2        N/A       MASTER       yes
```

➢ ***To configure Forced Master on a switch interface:***

Step 1. Enable Forced Master on the VLAN interface. Run:

```
switch (config) # interface vlan 2 ptp enable forced-master
```

Step 2. Enable Forced Master on the interface. Run:

```
switch (config) #  1/1 ptp enable forced-master
```

The interface should be a member in the PTP enabled VLAN(s).

➢ ***To verify PTP configuration:***

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Enabled
Domain                  : 127
Clock identity          : 7C:FE:90:FF:FE:FA:21:88
GMC identity            : 7C:FE:90:FF:FE:FA:21:88
Number of master ports  : 2
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------
  Port           VLAN      State        Forced Master
  ------------------------------------------------
  Eth1/1         2         MASTER       yes
  Eth1/1         3         SLAVE        no
```

Forced Master is indicated as “yes” only if enabled on the interface and the corresponding VLAN interface.

4.4.6 Commands

protocol ptp

protocol ptp

Enables PTP on the switch.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4110
Role	admin

Example	`switch (config) # protocol ptp`
	`...`
	`switch (config) #`

Related Commands	N/A
Notes	

ptp amt

ptp amt <clock-id>
no ptp amt <clock-id>

Adds an acceptable master table entry.
The no form of the command removes an acceptable master entry.

Syntax Description	interval	Range: -3 to 1
Default	N/A	
Configuration Mode	config	
History	3.6.8100	
Role	admin	
Example	`switch (config) # ptp amt 00:11:22:FF:FE:33:44:55:66`	
Related Commands	N/A	
Notes		

ptp announce interval

ptp announce interval <interval>

Configures PTP announce interval.

Syntax Description	interval	Range: -3 to 1
Default	-2	
Configuration Mode	config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
Role	admin	

Example	`switch (config  1/1) # ptp announce interval -2`
	`...`
	`switch (config  1/1) #`	

Related Commands	N/A	
Notes		

ptp announce timeout

ptp announce timeout <timeout>

Configures PTP announce timeout.

Syntax Description	timeout	Range: 2-10
Default	3	
Configuration Mode	config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
Role	admin	

Example	`switch (config  1/1) # ptp announce timeout 3`
	`...`
	`switch (config  1/1) #`	

Related Commands	N/A	
Notes		

ptp delay-req interval

ptp delay-req interval <interval>

Configures PTP delay-req interval.

Syntax Description	interval	Range is LogSyncInt -7 to 7
Default	-3	
Configuration Mode	config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
Role	admin	

Example	`switch (config  1/1) # ptp delay-req interval -3`
	`...`
	`switch (config  1/1) #`	

Related Commands	N/A	
Notes		

ptp domain

ptp domain <domain number>

Inserts the number of ptp domain.

Syntax Description	domain number	Range: 0-127
Default	127	
Configuration Mode	config	
History	3.6.4110	
Role	admin	

Example	`switch (config) # ptp domain`
	`...`
	`switch (config) #`	

Related Commands	N/A	
Notes		

ptp enable

	ptp enable Enables PTP per interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ethernet config interface vlan
History	3.6.4110 3.6.8008 Added “interface vlan” configuration mode
Role	admin

Example	`switch (config  1/1) # ptp enable`
	`...`
	`switch (config  1/1) #`

Related Commands	N/A
Notes	

ptp enable forced-master

	ptp enable forced-master **no ptp enable forced-master** Configures PTP interfaces to forced master state. The no form of the command removes PTP interfaces from forced master state.
Syntax Description	N/A
Default	no ptp enable forced-master
Configuration Mode	config interface ethernet config interface vlan
History	3.6.8100
Role	admin
Example	`switch (config  1/1) # ptp enable forced-master`
Related Commands	N/A
Notes	

ptp message-format

ptp message-format {mixed | multicast}

Configures PTP delay request messages format.

Syntax Description	mixed	Sends unicast delay request packets
	multicast	Sends multicast delay request packets
Default	mixed	
Configuration Mode	config	
History	3.6.8008	
Role	admin	
Example	`switch (config) # ptp message-format mixed`	
Related Commands	N/A	
Notes		

ptp priority

ptp priority{1 | 2} <priority>

Configures PTP primary priority.

Syntax Description	priority	Range: 0-255
Default	128	
Configuration Mode	config	
History	3.6.4110	
Role	admin	

Example	`switch (config) # ptp priority1 128`
	`...`
	`switch (config) #`	

Related Commands	N/A	
Notes		

ptp sync interval

ptp sync interval <interval>

Configures PTP sync interval.

Syntax Description	interval	Range: -1 to -7
Default	-3	
Configuration Mode	config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
Role	admin	

Example	`switch (config  1/1) # ptp sync interval -3`
	`...`
	`switch (config  1/1) #`	

Related Commands	N/A	
Notes		

clear ptp amt log

	clear ptp amt log Clears log of received clock IDs outside of acceptable master table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin
Example	`switch (config) # clear ptp amt log`
Related Commands	N/A
Notes	

clear ptp forced-master log

	clear ptp forced-master log Clears log of received clock IDs on forced master interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin
Example	`switch (config) # clear ptp forced-master log`
Related Commands	N/A
Notes	

clear ptp interface vlan ethernet counters

clear ptp interface vlan <id> ethernet <slot>/<port>[/<subport>] counters

Clears PTP counters for specified VLAN member interface.

Syntax Description	id	VLAN ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	
Example	`switch (config 1/1) # clear ptp interface vlan 2 ethernet 1/1 counters`	
Related Commands	N/A	
Notes		

clear ptp interface port-channel counters

clear ptp interface port-channel <id> counters

Clears PTP interface for LAG counters.

Syntax Description	id	LAG ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Role	Admin	
Example	`switch (config) # clear ptp interface port-channel 3 counters`	
Related Commands	N/A	
Notes		

clear ptp VRF counters

clear ptp vrf <vrf-name> counters

Clears the PTP VRF counters.

Syntax Description	vrf-name	Name of PTP enabled VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Role	Admin	
Example	`switch (config) # clear ptp vrf cust1 counters`	
Related Commands	N/A	
Notes	This command clears interface statistics on all PTP enabled interfaces in a specific PTP enabled VRF.	

interface port-channel

	interface port-channel * ptp enable [forced-master] **[no] interface port-channel * ptp enable [forced-master]** This command enables/disables PTP on port-channel Running the no form of this command will disable PTP on port-channel.	
Syntax Description	PTP is enabled on LAG interface	
	No-command attributes	[no] interface port-channel <id> ptp announce interval [<min> - <max>]
		[no] interface port-channel <id> ptp announce timeout [<min> - <max>]
		[no] interface port-channel <id> ptp delay-req interval [<min> - <max>]
		[no] interface port-channel <id> ptp sync interval [<min> - <max>]
Default	N/A	
Configuration Mode	Configure terminal	
History	3.7.1000	
Role	Admin	
Example	`switch (config) # interface port-channel 1 ptp enable forced-master`	
Related Commands	show ptp show ptp interface port channel * show ptp forced-master show ptp interface port-channel * counters clear ptp interface port-channel * counters [no] interface port-channel * ptp announce interval [(-3) - 1] [no] interface port-channel * ptp announce timeout [2 - 10] [no] interface port-channel * ptp delay-req interval [0 - 5] [no] interface port-channel * ptp sync interval [(-7) - (-1)]	
Notes	PTP need to be enabled on LAG members as well: interface ethernet * ptp enable [forced-master	

ptp vrf

	ptp vrf * enable [forced-master] **[no] ptp vrf * enable [forced-master]** This command enables/disables PTP in VRF. Running the no form of this command will disable PTP in <vrf-name> VRF.	
Syntax Description	PTP is enabled in VRF	
	No-command attributes	[no] ptp vrf <vrf-name> announce interval [<min> - <max>]
		[no] ptp vrf <vrf-name> announce timeout [<min> - <max>]
		(config)# [no] ptp vrf <vrf-name> delay-req interval [<min> - <max>]
		[no] ptp vrf <vrf-name> sync interval [<min> - <max>]
Default	N/A	
Configuration Mode	Configure terminal	
History	3.7.1000	
Role	Admin	
Example	`switch (config) # ptp vrf cust1 enable forced-master`	
Related Commands	show ptp show ptp vrf * show ptp forced-master show ptp vrf * counters clear ptp vrf * counters [no] ptp vrf * announce interval [(-3) - 1] [no] ptp vrf * announce timeout [2 - 10] [no] ptp vrf * delay-req interval [0 - 5] [no] ptp vrf * sync interval [(-7) - (-1)]	
Notes	PTP needs to be enabled on interfaces in VRF as well: interface ethernet * ptp enable [forced-master] interface port-channel * ptp enable [forced-master] interface vlan * ptp enable [forced-master]	

show ptp

	show ptp Displays PTP configuration and operation data.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110 3.6.8008 Updated Example 3.6.8100 Updated Example
Role	admin

Example	`switch (config) # show ptp`
	`PTP mode : Boundary Clock`
	`Message format : Mixed`
	`Acceptable Master Table : Disabled`
	`Domain : 127`
	`Clock identity : 7C:FE:90:FF:FE:FA:23:88`
	`GMC identity : 7C:FE:90:FF:FE:FA:23:88`
	`Number of master ports : 0`
	`Slave port interface : N/A`
	
	`PTP enabled interfaces:`
	`---------------------------------------------------------------------`
	`Port Po VLAN VRF State Forced Master`
	`---------------------------------------------------------------------`
	`Eth1/11 3 N/A default MASTER yes`
	`Eth1/10 3 N/A default MASTER yes`

Related Commands	N/A
Notes	

show ptp vrf

show ptp vrf <vrf_name>

Shows interfaces in VRF PTP related data.

Syntax Description	<vrf_name> - name of PTP enabled VRF
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
Role	Admin

Example	switch (config) # show ptp vrf Interface name: Eth1/1 Channel group ID: N/A VRF: cust1 IP Address: 1.1.1.1 Port Clock identity: E4:1D:2D:FF:FE:44:65:C8 PTP Port number: 1 PTP interface state: MASTER Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 Interface name: Eth1/2 Channel group ID: N/A VRF: default IP Address: 2.2.2.2 Port Clock identity: E4:1D:2D:FF:FE:44:65:C8 PTP Port number: 1 PTP interface state: SLAVE Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 Interface name: Eth1/1 Channel group ID: N/A VRF: cust1 IP Address: 1.1.1.1 Port Clock identity E4:1D:2D:FF:FE:44:65:C8 PTP Port number: 1 PTP interface state: MASTER Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4
Related Commands	N/A
Notes	Displays ptp state of all PTP-enabled interfaces in all PTP-enabled VRFs.

show ptp vrf counters

	show ptp vrf * counters Shows port statistics on interfaces in VRF.
Syntax Description	<vrf_name> - name of PTP enabled VRF
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
Role	Admin

Example

```
switch (config) # show ptp vrf cust1 counters
VRF: cust1

Eth1/1

RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count
0         Forwarded Management message count

Eth1/2

RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count
0         Forwarded Management message count
```

Related Commands	N/A
Notes	Display ptp counters of all PTP enabled interfaces in specific PTP enabled VRF.

show ptp amt

	show ptp amt Displays acceptable master table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin

Example	`switch (config) # show ptp amt`
	`Clock Identities:`
	`00:11:22:FF:FE:44:55:66`
	`66:55:44:FF:FE:22:11:00`

Related Commands	N/A
Notes	

show ptp interface port-channel

show ptp interface port-channel <po_id>

Shows LAG member interfaces PTP related data.

Syntax Description	po_id - Port channel (LAG) ID
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
Role	Admin

Example	`(config) # show ptp interface port-channel 3`
	`Interface name:  Eth1/10`
	`Channel group ID:  3`
	`VRF:  default`
	`IP Address:  1.1.1.2`
	`Port Clock identity:  7C:FE:90:FF:FE:FA:23:88`
	`PTP Port number:  3`
	`PTP interface state:  MASTER`
	`Forced Master:  no`
	`Delay request interval(log mean):  0`
	`Announce receipt time out:  3`
	`Announce interval(log mean):  0`
	`Sync interval(log mean):  -3`
	`Delay Mechanism:  End to End`
	`Transport protocol:  UDP IPv4`
	
	`Interface name:  Eth1/11 (Po 3)`
	`Channel group ID:  3`
	`VRF:  default`
	`IP Address:  1.1.1.2`
	`Port Clock identity:  7C:FE:90:FF:FE:FA:23:88`
	`PTP Port number:  2`
	`PTP interface state:  MASTER`
	`Forced Master:  no`
	`Delay request interval(log mean):  0`
	`Announce receipt time out:  3`
	`Announce interval(log mean):  0`

Related Commands	N/A
Notes	

show ptp interface port-channel counters

show ptp interface port-channel <po_id> counters

Shows port statistics on LAG member interfaces.

Syntax Description	po_id - Port channel (LAG) ID
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
Role	Admin

Example

```
(config) # show ptp interface port-channel 3 counters
Eth1/10
RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
1         Management message count
0         Forwarded Management message count

Eth1/11 (Po 3)
RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signalling message count
2         Management message count
0         Forwarded Management message count
```

Related Commands N/A

Notes

show ptp amt log

show ptp amt log

Displays received GMC clock IDs outside of acceptable master table.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin
Example	

```
switch (config) # show ptp amt log
-------------------------------------------------------------------------------
Clock Identity            Interface   VLAN   IP Address     Last Occurrence
-------------------------------------------------------------------------------
04:1D:2D:FF:FE:A5:F3:94   Eth1/2      N/A    192.168.66.7   2018/07/17 19:44:09
03:1D:2D:FF:FE:A5:F3:94   Eth1/2      N/A    192.168.66.7   2018/07/17 19:44:09
```

Related Commands	N/A
Notes	

show ptp clock

	show ptp clock Displays configuration and operation data of PTP clock.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Role	admin

Example	`switch (config) # show ptp clock`
	`Domain:  127`
	`Number of PTP ports:  1`
	`Priority1:  128`
	`Priority2:  128`
	`Clock identity:  e41d2d.fffe.46f801`
	`Offset From Master (ns): 65535`
	`Mean path delay (ns):  13303808`
	`Clock Quality`
	`Class:  248`
	`Accuracy:  254`
	`Offset (log variance): 65535`
	`Steps Removed from GMC:  1`
	`Local clock time:  13:59:27 Etc/UTC 2017/05/23`
	
	`...`

Related Commands	N/A
Notes	

show ptp clock parent

	show ptp clock parent Displays configuration and operation data of parent PTP clock.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Role	admin

Example	`switch (config) # show ptp clock parent`
	`Parent Clock`
	`Parent Clock identity:  7cfe90.fffe.fa2141`
	`Parent Port number:  2`
	
	`GMC`
	`GMC Identity:  7cfe90.fffe.fa2141`
	
	`GMC Clock Quality`
	`  Priority1:  128`
	`  Priority2:  128`
	`  Class:  248`
	`  Accuracy:  254`
	`  Offset (log variance):  65535`
	
	`...`

Related Commands	N/A
Notes	

show ptp forced-master

show ptp forced-master

Displays forced master PTP interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin

Example	`(config) # show ptp forced-master`
	`---------------------------------------------`
	`Port  Po  VLAN  VRF`
	`---------------------------------------------`
	`Eth1/10  3  N/A  default`
	`Eth1/11  3  N/A  default`

Related Commands	N/A
Notes	

show ptp forced-master log

show ptp forced-master log

Displays clock IDs received on forced master interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin
Example	

```
switch (config) # show ptp forced-master log
------------------------------------------------------------------------
Clock Identity           Interface  VLAN  IP Address    Last Occurrence
------------------------------------------------------------------------
04:1D:2D:FF:FE:A5:F3:94  Eth1/2     N/A   192.168.66.7  2018/07/17 19:44:09
03:1D:2D:FF:FE:A5:F3:94  Eth1/2     N/A   192.168.66.7  2018/07/17 19:44:09
```

Related Commands	N/A
Notes	

show ptp

show ptp <slot>/<port>[/<subport>]

Displays PTP configuration and operation data per Ethernet port.

Syntax Description	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8100	Updated Example
Role	admin	

Example	`switch (config) # show ptp 1/1`
	`Interface name: Eth1/1`
	`Port Clock identity: 7cfe90.fffe.fa21c1`
	`PTP Port number: 1`
	`PTP interface state: SLAVE`
	`Forced Master: no`
	`Delay request interval(log mean): -3`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`
	
	`...`	

Related Commands	N/A	
Notes		

show ptp counters

show ptp <slot>/<port>[/<subport>] counters

Displays PTP counters per Ethernet port.

Syntax Description	<slot>/<port>/<subport> Ethernet port ID (e.g. 1/3/1)
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110 3.6.8008 Added VLAN parameter
Role	admin

Example	`switch (config) # show ptp 1/5 counters`
	`Eth1/5`
	`RX`
	`108 Sync message count`
	`0 Delay request message count`
	`0 PDelay request message count`
	`0 PDelay response message count`
	`108 Follow Up message count`
	`17 Delay response message count`
	`0 PDelay response follow Up message count`
	`54 Announce message count`
	`0 Signaling message count`
	`0 Management message count`
	
	`TX`
	`74188 Sync message count`
	`17 Delay request message count`
	`0 PDelay request message count`
	`0 PDelay response message count`
	`74188 Follow Up message count`
	`0 Delay response message count`
	`0 PDelay response follow Up message count`
	`37117 Announce message count`
	`0 Signaling message count`
	`57 Management message count`
	
	`...`

Related Commands	N/A
Notes	

show ptp interface vlan

show ptp interface vlan <vid>

Displays PTP configuration and operation data per VLAN.

Syntax Description	vid	VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
	3.6.8100	Updated Example
Role	admin	

Example	`switch (config) # show ptp interface vlan 1`
	`Interface name: Eth1/15/1 (VLAN 1)`
	`Port Clock identity: 7cfe90.fffe.fa2388`
	`PTP Port number: 1`
	`PTP interface state: SLAVE`
	`Forced Master: no`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`
	
	`...`	

Related Commands	N/A	
Notes		

show ptp interface vlan ethernet

show ptp interface vlan <vid> ethernet <slot>/<port>[/<subport>]

Displays PTP configuration and operation data for specified VLAN member interface for a specified Ethernet port.

Syntax Description	vid	VLAN ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ptp interface vlan 1 ethernet 1/15/1`
	`Interface name: Eth1/15/1 (VLAN 1)`
	`Port Clock identity: 7cfe90.fffe.fa2388`
	`PTP Port number: 1`
	`PTP interface state: FAULTY`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`	

Related Commands	N/A	
Notes		

show ptp interface vlan counters

show ptp interface vlan <vid> counters

Displays PTP counters per VLAN.

Syntax Description	vid	VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example

```
switch (config) # show ptp interface vlan 1 counters
Eth1/15/1 (VLAN 1)
RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signaling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signaling message count
0         Management message count

Eth1/15/2 (VLAN 1)
RX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signaling message count
0         Management message count

TX
0         Sync message count
0         Delay request message count
0         PDelay request message count
0         PDelay response message count
0         Follow Up message count
0         Delay response message count
0         PDelay response follow Up message count
0         Announce message count
0         Signaling message count
0         Management message count
```

Related Commands N/A

Notes

show ptp interface vlan ethernet counters

show ptp interface vlan <vid> ethernet <slot>/<port>[/<subport>] counters

Displays PTP counters per VLAN for a specified Ethernet port.

Syntax Description	vid	VLAN ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ptp interface vlan 1 ethernet 1/15/1 counters`
	`Eth1/15/1 (VLAN 1)`
	`RX`
	`0 Sync message count`
	`0 Delay request message count`
	`0 PDelay request message count`
	`0 PDelay response message count`
	`0 Follow Up message count`
	`0 Delay response message count`
	`0 PDelay response follow Up message count`
	`0 Announce message count`
	`0 Signaling message count`
	`0 Management message count`
	
	`TX`
	`0 Sync message count`
	`0 Delay request message count`
	`0 PDelay request message count`
	`0 PDelay response message count`
	`0 Follow Up message count`
	`0 Delay response message count`
	`0 PDelay response follow Up message count`
	`0 Announce message count`
	`0 Signaling message count`
	`0 Management message count`	

Related Commands	N/A	
Notes		

4.5 Software Management

4.5.1 Important Pre-OS Upgrade Notes

Please consider the following items prior to upgrading the OS:

- The system becomes unavailable while OS upgrade is running
- The upgrade procedure burns the software image as well as the firmware should there be a need
- Before upgrading the software image on your system, make sure to close all CLI sessions besides the one used to run the upgrade process
- To upgrade the Mellanox Onyx™ version on an MLAG cluster, please refer to Section 4.5.3, “Upgrading Onyx HA Groups,” on page 278
- Interfaces with global pause are not mapped to a lossless pool after upgrade from versions earlier than 3.6.5000
- You have to read and accept the End-User License Agreement (EULA) after image upgrade in case the EULA is modified. The EULA link is only available upon first login to CLI.
- Linux docker container names are limited to 180 characters. Upgrading to this version removes containers which do not comply with this limitation and prints the following warning to the log: “Removed configuration of container: <container name>, container name is limited to 180 characters”.
- When upgrading from a version older than 3.6.3130 with an MLAG cluster, output appears as in “UP” and “Peering” state instead of “Upgrade” on both MLAG VIP clusters. The upgrade process will not be affected.

4.5.2 Upgrading Onyx Software

➢ ***To upgrade Onyx on your system, perform the following steps:***

Step 1. Enter Config mode. Run:

```
switch > enable
switch # configure terminal
switch (config) #
```

Step 2. Display the currently available image (.img file).

```
switch (config) # show images
Installed images:

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <old_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

Step 3. Delete the image listed under “Images available to be installed” prior to fetching the new image. Use the command “image delete” for this purpose.

```
switch (config) # image delete <old_image>
```

When deleting an image, you delete the file but not the partition. This is recommended so as to not overload system resources.

Step 4. Fetch the new software image.

```
switch (config) # image fetch scp://<username>:<password>@<ip-address>/var/www/html/
<new_image>
Password (if required): ******
100.0%[##########################################################s###############]
```

Step 5. Display the available images again and verify that the new image now appears under “Images available to be installed”.

To recover from image corruption (e.g. due to power interruption), there are two installed images on the system. See the commands “image boot next”, and “image boot location” for more information.

```
switch (config) # show images
Installed images:

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <new_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

Step 6. Install the new image.

```
switch (config) # image install <new_image>
Step 1 of 4: Verify Image
 100.0%  [################################################################]
Step 2 of 4: Uncompress Image
 100.0%  [################################################################]
Step 3 of 4: Create Filesystems
 100.0%  [################################################################]
Step 4 of 4: Extract Image
 100.0%  [################################################################]
```

CPU utilization may go up to 100% during image upgrade.

Step 7. Have the new image activate during the next boot. Run:

```
switch (config) # image boot next
```

Step 8. Run “show images” to review your images. Run:

```
switch (config) # show images
Installed images:

  Partition 1:
  <new_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <new_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

Step 9. Save current configuration. Run:

```
switch (config) # configuration write
```

Step 10. Reboot the switch to run the new image. Run:

```
switch (config) # reload
Configuration has been modified; save first? [yes] yes
Configuration changes saved.
Rebooting...
switch (config)#
```

After software reboot, the software upgrade will also automatically upgrade the firmware version.

When performing upgrade from the WebUI, make sure that the image you are trying to upgrade to is not located already in the system (i.e. fetched from the CLI).

4.5.3 Upgrading Onyx HA Groups

In case fallback is ever necessary in an HA group, all cluster nodes must have the same Onyx version installed and they must be immediately reloaded.

➢ ***To upgrade Onyx version without affecting an HA group:***

Step 1. Identify the HA group master.

for MLAG. Run:

```
switch (config)# show mlag-vip
MLAG VIP
========
MLAG group name: my-mlag-group
MLAG VIP address: 1.1.1.1/30
Active nodes: 2

Hostname            VIP-State           IP Address
--------------------------------------------------
SwitchA             master              10.10.10.1
SwitchB             standby             10.10.10.2
```

Step 2. Upgrade standby node in the HA group according to steps 1-10 in Section •, “When upgrading from a version older than 3.6.3130 with an MLAG cluster, output appears as in “UP” and “Peering” state instead of “Upgrade” on both MLAG VIP clusters. The upgrade process will not be affected.,” on page 274.

Step 3. Wait until all standby nodes have rejoined the group.

In situations of heavy CPU load or noisy network, it is possible that another node assumes the role of cluster master before all standby nodes have rejoined the group. If this happens, you may stop waiting and proceed directly to Step 4.

Step 4. Upgrade the master node in the HA group according to steps 1-10 in Section •, “When upgrading from a version older than 3.6.3130 with an MLAG cluster, output appears as in “UP” and “Peering” state instead of “Upgrade” on both MLAG VIP clusters. The upgrade process will not be affected.,” on page 274.

4.5.4 Upgrading Onyx MLAG-STP Setup

To upgrade the OS on an MLAG-STP setup from 3.6.610x to this version, there are two possible procedures:

Procedure 1:

Step 1. Make sure there are no loops in the fabric.

Step 2. Disable STP. Run:

```
switch (config) # no spanning-tree
```

Step 3. Perform the upgrade according to steps 1-10 in Section 4.5.3, on page 278.

Step 4. Enable STP – this step may lead to traffic loss while the STP state is converging. Run:

```
switch (config) # spanning-tree
```

Procedure 2:

Step 1. Shutdown all ports on the MLAG slave.

Step 2. Save configuration. Run:

```
switch (config) # configuration write
```

Step 3. Upgrade MLAG slave according to steps 1-10 in Section 4.5.3, on page 278.

Step 4. Upgrade MLAG master. Run:

```
switch (config) # reload force immediate
```

Step 5. Enable all ports on the MLAG slave.

4.5.5 Deleting Unused Images

➢ ***To delete unused images:***

Step 1. Enter Config mode. Run:

```
switch >
switch > enable
switch # configure terminal
```

Step 2. Get a list of the unused images. Run

```
switch (config) # show images
Images available to be installed:
  image-PPC_M460EX-3.1.1224.img
  SX-OS_PPC_M460EX 3.1.1224 2011-04-28 12:29:48 ppc
Installed images:
Partition 1:
SX-OS_PPC_M460EX 3.1.0000-dev-HA 2011-04-10 12:02:49 ppc
Partition 2:
SX-OS_PPC_M460EX 3.1.0000-dev-HA 2011-04-10 12:02:49 ppc

Last boot partition: 1
Next boot partition: 1
Boot manager password is set.
No image install currently in progress.
Require trusted signature in image being installed: yes
switch (config) #
```

Step 3. Delete the unused images. Run:

```
switch config) # image delete image-X86_64-3.6.3234-12.img
switch (config) #
```

When deleting an image, you delete the file but not the partition. This is recommended so as to not overload system resources.

4.5.6 Downgrading Onyx Software

Prior to downgrading software, please make sure the following prerequisites are met:

Step 1. Log into your switch via the CLI using the console port.

Step 2. Backup your configuration according to the following steps:

1. Change to Config mode. Run:

```
switch-112094 [standalone: master] > enable
switch-112094 [standalone: master] # configure terminal
switch-112094 [standalone: master] (config) #
```

2. Disable paging of CLI output. Run:

```
switch-112094 [standalone: master] (config) # no cli default paging enable
```

3. Display commands to recreate current running configuration. Run:

```
switch-112094 [standalone: master] (config) # show running-config
```

4. Copy the output to a text file.

4.5.6.1 Downloading Image

Step 1. Log into your system to obtain its product number. Run:

```
switch-112094 [standalone: master] (config) # show inventory
```

Step 2. Log into MyMellanox at https://mymellanox.force.com/support/SupportLoginand download the relevant Onyx version to your system type.

Step 3. Log into the switch via the CLI using the console port.

Step 4. Change to Config mode. Run:

```
switch > enable
switch # configure terminal
switch (config) #
```

Step 5. Delete all previous images from the `Images available to be installed` prior to fetching the new image. Run:

```
switch (config) # image-X86_64-3.6.3234-12.img
```

Step 6. Fetch the requested software image. Run:

```
switch (config) # image fetch scp://username:password@192.168.10.125/var/www/html/
<image_name>
100.0%[################################################## ###############]
```

4.5.6.2 Downgrading Image

The procedure below assumes that booting and running is done from Partition 1 and the downgrade procedure is performed on Partition 2.

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Display all image files on the system. Run:

```
switch (config) # show images
Images available to be installed:
new_image.img
<downgrade version> 2010-09-19 16:52:50
Installed images:
Partition 1:
<current version> 2010-09-19 03:46:25
Partition 2:
<current version> 2010-09-19 03:46:25
Last boot partition: 1
Next boot partition: 1
No boot manager password is set.
switch (config) #
```

Step 4. Install the Onyx image. Run:

```
switch (config) # image install <image_name>
Step 1 of 4: Verify Image
100.0% [########################################################################]
Step 2 of 4: Uncompress Image
100.0% [########################################################################]
Step 3 of 4: Create Filesystems
100.0% [########################################################################]
Step 4 of 4: Extract Image
100.0% [########################################################################]
switch (config) #
```

Step 5. Display all image files on the system. Run:

```
switch (config) # show images
Images available to be installed:
new_image.img
 <downgrade version> 2010-09-19 16:52:50
Installed images:
Partition 1:
 <current version> 2010-09-19 03:46:25
```

```
Partition 2:
 <downgrade version> 2010-09-19 16:52:50
Last boot partition: 1
Next boot partition: 2
No boot manager password is set.
switch (config) #
```

Step 6. Configure the boot location to be the other (next) partition. Run:

```
switch (config) # image boot next
```

There are two installed images on the system. Therefore, if one of the images gets corrupted (due to power interruption, for example), in the next reboot the image will go up from the second partition.

In case you are downloading to an older software version which has never been run yet on the switch, use the following command sequence as well:

switch (config) # no boot next fallback-reboot enable

switch (config) # configuration write

Step 7. Reload the switch. Run:

```
switch (config) # reload
```

4.5.6.3 Switching to Partition with Older Software Version

The system saves a backup configuration file when upgrading from an older software version to a newer one. If the system returns to the older software partition, it uses this backup configuration file.

*****IMPORTANT NOTE*****

All configuration changes done with the new software are lost when returning to the older software version.

There are 2 instances where the backup configuration file does not exist:

- The user has run “reset factory” command, which clears all configuration files in the system
- The user has run “configuration switch-to” to a configuration file with different name then the backup file

Note that the configuration file becomes empty if the switch is downgraded to a software version which has never been installed yet.

To allow switching partition to the older software version for the 2 aforementioned cases only, follow the steps below:

Step 1. Run the command:

```
switch (config)# no boot next fallback-reboot enable
```

Step 2. Set the boot partition. Run:

```
switch (config)# image boot next
```

Step 3. Save the configuration. Run:

```
switch (config)# configuration write
```

Step 4. Reload the system. Run:

```
switch (config)# reload
```

4.5.7 Upgrading System Firmware

Each Onyx software package version has a default switch firmware version. When you update the Onyx software to a new version, an automatic firmware update process will be attempted by Onyx. This process is described below.

4.5.7.1 After Updating Onyx Software

Upon rebooting your switch system after updating the Onyx software, Onyx compares its default firmware version with the currently programmed firmware versions on all the switch modules (leafs and spines on director-class switches, or simply the switch card on edge switch systems).

If one or more of the switch modules is programmed with a firmware version other than the default version, then Onyx automatically attempts to burn the default firmware version instead.

If a firmware update takes place, then the login process is delayed a few minutes.

To verify that the firmware update was successful, log into Onyx and run the command “show asic-version” (can be run in any mode). This command lists all of the switch modules along with their firmware versions. Make sure that all the firmware versions are the same and match the default firmware version. If the firmware update failed for one or more modules, then the following warning is displayed.

Some subsystems are not updated with a default firmware.

If you detect a mismatch in firmware version for one or more modules of the switch system, please contact your assigned Mellanox Technologies field application engineer.

4.5.7.2 Importing Firmware and Changing the Default Firmware

To perform an automatic firmware update by Onyx for a different switch firmware version without changing the Onyx version, import the firmware package as described below. Onyx sets it as the new default firmware and performs the firmware update automatically as described in the previous subsections.

4.5.7.2.1 Default Firmware Change on Standalone Systems

Step 1. Import the firmware image (.mfa file). Run:

```
switch (config) # image fetch scp://root@1.1.1.1:/tmp/fw-SIB-rel-11_1600_0200-FIT.mfa
Password (if required): *******
100.0%
[##############################################################################]
switch (config) # image default-chip-fw fw-SIB-rel-11_1600_0200-FIT.mfa
Installing default firmware image. Please wait...
Default Firmware 11.1600.0200 updated. Please save configuration and reboot for new FW
to take effect.
```

Step 2. Save the configuration. Run:

```
switch (config) # configuration write
```

Step 3. Reboot the system to enable auto update.

4.5.8 Image Maintenance via Mellanox ONIE

ONIE is an “open compute” Open Network Install Environment for bare metal network switches. ONIE enables a bare metal network switch ecosystem where end-users have a choice among different network operating systems.

Onyx is distributed in way that allows installation on an ONIE environment. Certain Mellanox switch models come pre-installed with ONIE and Onyx and support changing to a different operating system (OS).

➢ ***To change the switch system’s OS:***

Step 1. Reboot the switch and wait for it to reach the GRUB menu:

```
                                GNU GRUB version 2.02

X86_64 3.4.1932 2015-04-24 18:04:12 x86_64 1
X86_64 3.4.1932 2015-04-24 18:04:12 x86_64 2
ONIE
```

Step 2. Select the ONIE option using the arrow keys. The following message appears:

```
Due to security constraints, this option will uninstall your current MLNX OS system.
Are you sure ?
```

Step 3. Type YES to continue.

Since Onyx is being uninstalled and deleted from the hard drive, the process takes a few hours. After this is finished, the system reboots into the ONIE shell and auto discovery begins.

```
Info: Fetching tftp://<ip-address>/7C-FE-90-5E-6A-4A/onie-installer-x86_64-mlnx_x86-
r5.0.1400 ...
Failure: Unable to find installer: /installer
Info: Fetching tftp://<ip-address>/0AE016FB/onie-installer-x86_64-mlnx_x86-r5.0.1400
...
Failure: Unable to find installer: /installer
Info: Fetching tftp://<ip-address>/0AE016F/onie-installer-x86_64-mlnx_x86-r5.0.1400
...
...
```

Step 4. In order to manually insert an install URL, press Enter and insert the command “install_url <http> / <tftp> <url> <image name .bin>”. For example:

```
install_url http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-3.5.1000-21-
installer.bin
```

Once you hit Enter, you have about 4 second to insert the command so it is recommended to prepare the command in advance and simply pasting it in. At this stage, the OS installation begins.

Step 5. Wait for the installation to end and reboot this switch to boot into the OS.

```
ONIE:/ # install_url http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_
64/X86_64-3.5.1000-21-installer.bin
Stopping: discover... done.
down.
ONIE: eth1: link down.  Skipping configuration.
ONIE: Failed to configure eth1 interface
Info: Fetching http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-3.5.1000-21-
installer.bin ...
Connecting to <ip_address>
installer           100% |*******************************|   392M  0:00:00 ETA
ONIE: Executing installer: http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-
3.5.1000-21-installer.bin
```

4.5.9 Commands

image boot

image boot {location <location ID> | next}

Specifies the default location where the system should be booted from.

Syntax Description	location ID	Specifies the default destination location. There can be up to 2 images on the system. The possible values are 1 or 2.
	next	Sets the boot location to be the next once after the one currently booted from, thus avoiding a cycle through all the available locations.
Default	N/A	
Configuration Mode	enable/config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # image boot location 2` `switch (config) #`	
Related Commands	show images	
Notes		

boot next

boot next fallback-reboot enable
no boot next fallback-reboot enable

Sets the default setting for next boot. Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate), it will reboot to the other partition as a fallback.
The no form of the command tells the system not to do that, only for the next boot.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.0506
Role	admin
Example	`switch (config) # boot next fallback-reboot enable` `switch (config) #`

Related Commands	show images
Notes	• Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate) it reboots to the other partition as a fallback. • The no form of this command tells the system not to do that **only** for the next boot. In other words, this setting is not persistent, and goes back to enabled automatically after each boot. • When downgrading to an older software version which has never been run yet on a system, the "fallback reboot" **always** happens, unless the command "no boot next fallback-reboot enable" is used. However, this also happens when the older software version *has* been run before, but the configuration file has been switched since upgrading. In general, a downgrade only works (without having the fallback reboot forcibly disabled) if the process can find a snapshot of the configuration file (by the same name as the currently active one) which was taken before upgrading from the older software version. If that is not found, a fallback reboot is performed in preference to falling back to the initial database because the latter generally involves a loss of network connectivity, and avoiding that is of paramount importance.

boot system

boot system {location | next}
no boot system next

Configures which system image to boot by default.
The no form of the command resets the next boot location to the current active one.

Syntax Description	location	Specifies location from which to boot system • 1 – installs to location 1 • 2 – installs to location 2
	next	Boots system from next location after one currently booted
Default	N/A	
Configuration Mode	config	
History	3.2.0506	
Role	admin	
Example	`switch (config) # boot system location 2` `switch (config) #`	
Related Commands	show images	
Notes		

image default-chip-fw

image default-chip-fw <filename>
no image default-chip-fw <original-fw-filename>

Sets the default firmware package to be installed.
The no form of the command resets default firmware package.

Syntax Description	filename	Specifies the firmware filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.6.6000	Added no form of the command
Role	admin	
Example	`switch (config) # image default-chip-fw fw-SPC-rel-13_1600_0184-FIT.mfa`	
Related Commands	show asic-version show images	
Notes		

image delete

image delete <image name>

Deletes the specified image file.

Syntax Description	image name	Specifies the image name.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # image delete image-MLXNX-OS-201140526-010145.img` `switch (config) #`	
Related Commands	show images	
Notes		

image fetch

image fetch <URL> [<filename>]

Downloads an image from the specified URL or via SCP.

Syntax Description	URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
	filename	Specifies a filename for this image to be stored as locally.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # image fetch scp://<username>@192.168.10.125/var/www/
html/<image_name>
Password  ******
100.0%[##################################################################]
switch (config) #

Other options:

switch (config) # image fetch http://10.1.0.40/path/filename
switch (config) # image fetch http://[fd4f:13:cc00:1::40]/path/filename
switch (config) # image fetch ftp://user:mypassword@10.1.0.40/foo/
bar.img
switch (config) # image fetch ftp://user:mypassword@[fd4f:13:cc00:1:
:40]/foo/bar.img
switch (config) # image fetch tftp://hostname/dir/filename
switch (config) # image fetch tftp://[fd4f:13:cc00:1::40]/dir/filename
switch (config) # image fetch scp://user@myhost/dir/filename
switch (config) # image fetch scp://user@myhost:1022/dir/filename
switch (config) # image fetch scp://user:pass@[fd4f:13:cc00:1::40]/dir/
filename
switch (config) # image fetch sftp://user@myhost/dir/filename
switch (config) # image fetch sftp://user@[fd4f:13:cc00:1::40]:1022/
dir/filename
switch (config) # image fetch sftp://user:pass@[fd4f:13:cc00:1::40]/
dir/filename
```

Related Commands	show images
Notes	• Please delete the previously available image, prior to fetching the new image • The path to the file in the case of TFTP depends on the server configuration. Therefore, it may not be an absolute path but a relative one. • See Section •, “When upgrading from a version older than 3.6.3130 with an MLAG cluster, output appears as in “UP” and “Peering” state instead of “Upgrade” on both MLAG VIP clusters. The upgrade process will not be affected.,” on page 274

image install

image install <image filename> [location <location ID>] | [progress <prog-options>]

Installs the specified image file.

Syntax Description	image filename	Specifies the image name.
	location ID	Specifies the image destination location.
	prog-options	• "no-track" overrides CLI default and does not track the installation progress • "track" overrides CLI default and tracks the installation progress
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # image install X86_64 4100-12 2017-07-26 06:54:12 x86_64`
	`Step 1 of 4: Verify Image`
	`100.0%`
	`[####################################################################]`
	`Step 2 of 4: Uncompress Image`
	`100.0%`
	`[####################################################################]`
	`Step 3 of 4: Create Filesystems`
	`100.0%`
	`[####################################################################]`
	`Step 4 of 4: Extract Image`
	`100.0%`
	`[####################################################################]`
	`switch (config) #`	

Related Commands	show images	
Notes	• The image cannot be installed on the "active" location (the one which is currently being booted) • On a two-location system, the location is chosen automatically if no location is specified	

image move

image move <src image name> <dest image name>

Renames the specified image file.

Syntax Description	src image name	Specifies the old image name.
	dest image name	Specifies the new image name.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # image move image1.img image2.img` `switch (config) #`	
Related Commands	show images	
Notes		

image options

image options serve all
no image options serve all

Configures options and defaults for image usage.
The no form of the command disables options and defaults for image usage.

Syntax Description	serve all	Specifies that the image files present on this appliance should be made available for HTTP and/or HTTPS download
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # image options serve all`	
Related Commands	show images	
Notes	The parameter “serve all” affects not only the files currently present, but also any files that are later downloaded. It only applies to image files, not the installed images, which are not themselves in a downloadable format. After running “serve all” the URLs where the images will be available are: • http://<HOSTNAME>/system_images/<FILENAME> • https://<HOSTNAME>/system_images/<FILENAME>	

show bootvar

show bootvar

Displays the installed system images and the boot parameters.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch [standalone: master] (config) # show bootvar`
	`Installed images:`
	
	`  Partition 1:`
	`  X86_64 3.6.4110-12 2017-07-26 06:54:12 x86_64`
	
	`  Partition 2:`
	`  X86_64 3.6.4006 2017-07-03 16:17:39 x86_64`
	
	`Last boot partition: 1`
	`Next boot partition: 1`
	
	`Serve image files via HTTP/HTTPS: no`
	
	`Boot manager password is set.`
	
	`Image signing: trusted signature always required`
	`Admin require signed images: yes`
	
	`Settings for next boot only:`
	`   Fallback reboot on configuration failure: yes (default)`
	`switch [standalone: master] (config) #`

Related Commands	N/A
Notes	

show images

	show image Displays information about the system images and boot parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch [standalone: master] (config) # show images`
	`Installed images:`
	
	`  Partition 1:`
	`  X86_64 3.6.4110-12 2017-07-26 06:54:12 x86_64`
	
	`  Partition 2:`
	`  X86_64 3.6.4006 2017-07-03 16:17:39 x86_64`
	
	`Last boot partition: 1`
	`Next boot partition: 1`
	
	`Images available to be installed:`
	
	`  webimage.tbz`
	`  X86_64 3.6.4071-12 2017-07-26 06:54:12 x86_64`
	
	`Serve image files via HTTP/HTTPS: no`
	
	`No image install currently in progress.`
	
	`Boot manager password is set.`
	
	`Image signing: trusted signature always required`
	`Admin require signed images: yes`
	
	`Settings for next boot only:`
	`   Fallback reboot on configuration failure: yes (default)`

Related Commands	N/A
Notes	

4.6 Configuration Management

4.6.1 Saving a Configuration File

To save the current configuration to the active configuration file, you can either use the `configuration write` command (requires running in Config mode) or the `write memory` command (requires running in Enable mode).

- To save the configuration to the active configuration file, run:

```
switch (config) # configuration write
```

- To save the configuration to a user-specified file without making the new file the active configuration file, run:

```
switch (config) # configuration write to myconf no-switch
```

- To save the configuration to a user-specified file and make the new file the active configuration file, run:

```
switch (config) # configuration write to myconf
```

- To display the available configuration files and the active file, run:

```
switch (config) # show configuration files
initial
myconf (active)
switch (config) #
```

4.6.2 Loading a Configuration File

By default, or after a system reset, the system loads the default "initial" configuration file.

➢ ***To load a different configuration file and make it the active configuration:***

```
switch [standalone: master] >
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
switch [standalone: master] (config) # configuration switch-to myconfig
switch [standalone: master] (config) #
```

4.6.3 Restoring Factory Default Configuration

If system configuration becomes corrupted, it is suggested to restore factory default configuration.

➢ ***To restore factory default configuration on a single management module system, run:***

```
switch (config) # reset factory keep-basic
```

4.6.4 Managing Configuration Files

There are two types of configuration files that can be applied on the switch, BIN files (binary) and text-based configuration files.

4.6.4.1 BIN Configuration Files

BIN configuration files are not human readable. Additionally, these files are encrypted and contain integrity verification preventing them from being edited and used on the switch.

- ***To create a new BIN configuration file:***

```
switch (config) # configuration new my-filename
```

A newly created BIN configuration file is always empty and is not created from the running-config.

- ***To upload a BIN configuration file from a switch to an external file server***:

```
switch (config) # configuration upload my-filename scp://myusername@my-server/path/to/
my/<file>
```

- ***To fetch a BIN configuration file***:

```
switch (config) # configuration fetch scp://myusername@my-server/path/to/my/<file>
```

- ***To see the available configuration files***:

```
switch (config) # show configuration files
initial (active)
my-filename

Active configuration: initial
Unsaved changes:      no
switch (config) #
```

- ***To load a BIN configuration file:***

```
switch (config) # configuration switch-to my-filename
This requires a reboot.
Type 'yes' to confirm: yes
```

Applying a new BIN configuration file changes the whole switch's configuration and requires system reboot which can be preformed using the command `reload`.

A binary configuration file uploaded from the switch is encrypted and has integrity verification. If the file is modified in any manner, the fetch to the switch fails.

4.6.4.2 Text Configuration Files

Text configuration files are text based and editable. It is similar in form to the output of the command "show running-config expanded".

- ***To create a new text-based configuration file***:

```
switch (config) # configuration text generate active running save my-filename
```

A newly created text configuration file is always created from the running-config.

- ***To apply a text-based configuration file:***

```
switch (config) # configuration text file my-filename apply
```

Applying a text-based configuration file to an existing/running data port configuration may result in unpredictable behavior. It is therefore suggested to first clear the switch's configuration by applying a specific configuration file (following the procedure in Section 4.6.4.1) or by resetting the switch back to factory default.

- ***To upload a text-based configuration file from a switch to an external file server***

```
switch (config) # configuration text file my-filename upload scp://root@my-server/root/
tmp/my-filename
```

- ***To fetch a text-based configuration file from an external file server to a switch***

```
switch (config) # configuration text fetch scp://root@my-server/root/tmp/my-filename
```

- ***To apply a text-based configuration file:***

```
switch (config) # configuration text file my-filename apply
```

When applying a text-based configuration file, the configuration is appended to the switch's existing configuration. Only new or changed configuration is added. Reboot is not required.

4.6.5 Commands

4.6.5.1 File System

debug generate dump

	debug generate dump Generates a debug dump.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # debug generate dump`
	`Generated dump sysdump-switch-112104-201140526-091707.tgz`
	`switch (config) #`

Related Commands	file debug-dump
Notes	The dump can then be manipulated using the “file debug-dump...” commands.

file debug-dump

file debug-dump {delete {<filename> | all | latest} | email {<filename> | latest} | upload {<filename> | latest} <URL>}

Manipulates debug dump files.

Syntax Description	delete	Deletes a debug dump file. • all: Deletes all existing debug files from this machine • latest: Deletes latest debug file from this machine
	email	Emails a debug dump file to pre-configured recipients for “informational events”, regardless of whether they have requested to receive “detailed” notifications or not. • latest: Emails the latest debug file to a pre-configured recipients
	upload	Uploads a debug dump file to a remote host. • latest: Uploads the latest debug file to a remote host
	URL	The URL to the remote host: HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4000	Added “all” and “latest” options
Role	admin	
Example	`switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz`	
Related Commands	show files debug-dump	
Notes		

file debug-dump

file debug-dump {delete {<filename> | latest} | email {<filename> | latest} | upload {{<filename> | latest} <URL>}}

Manipulates debug dump files.

Syntax Description	delete {<filename> | latest}	Deletes a debug dump file.
	email {<filename> | latest}	Emails a debug dump file to pre-configured recipients for “informational events”, regardless of whether they have requested to receive “detailed” notifications or not.
	upload {{<filename> | latest} <URL>}}	Uploads a debug dump file to a remote host. The URL to the remote host: HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4000	Added “latest” parameter
Role	admin	
Example	`switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz` `switch (config) #`	
Related Commands	show files debug-dump	
Notes		

file stats

file stats {delete <filename> | move {<source filename> | <destination filename>} | upload <filename> <URL>}

Manipulates statistics report files.

Syntax Description	delete <filename>	Deletes a stats report file.
	move <source filename> <destination filename>	Renames a stats report file.
	upload <filename> <URL>	Uploads a stats report file. URL - HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@host-name/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # file stats move memory-1.csv memory-2.csv` `switch (config) #`	
Related Commands	show files stats show files stats <filename>	
Notes		

file tcpdump

file tcpdump {delete <filename> | upload <filename> <URL>}

Manipulates tcpdump output files.

Syntax Description	delete <filename>	Deletes the specified tcpdump output file.
	upload <filename> <URL>	Uploads the specified tcpdump output file to the specified URL. URL - HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # file tcmpdump delete my-tcpdump-file.txt` `switch (config) #`	
Related Commands	show files stats tcpdump	
Notes		

reload

reload [force immediate | halt [noconfirm] | noconfirm]

Reboots or shuts down the system.

Syntax Description	force immediate	Forces an immediate reboot of the system even if the system is busy.
	halt	Shuts down the system.
	noconfirm	Reboots the system without asking about unsaved changes.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # reload`
	`Configuration has been modified; save first? [yes] yes`
	`Configuration changes saved.`
	`...`
	`switch (config) #`	

Related Commands	reset factory	
Notes		

reset factory

reset factory [keep-all-config | keep-basic | keep-virt-vols | only-config] [halt]

Clears the system and resets it entirely to its factory state.

Syntax Description	keep-all-cofig	Preserves all configuration files including licenses. Removes the logs, stats, images, snapshots, history, known hosts. The user is prompted for confirmation before honoring this command, unless confirmation is disabled with the command: “no cli default prompt confirm-reset”.
	keep-basic	Preserves licenses in the running configuration file
	keep-virt-vols	Preserve all virtual disk volumes
	only-config	Removes configuration files only. The logs, stats, images, snapshots, history, and known hosts are preserved.
	halt	The system is halted after this process completes
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added notes and “keep-virt-vols” parameter
	3.6.2002	Updated Example and Notes
Role	admin	

Example	`switch (config) # reset factory`
	`Warning - confirming will cause system reboot.`
	`Type 'YES' to confirm reset: YES`
	`Resetting and rebooting the system -- please wait...`
	`...`	

Related Commands	reload	
Notes	• Effects of parameter “keep-all-cofig”: Licenses – not deleted; profile – no change; configuration – unchanged; management IP – unchanged • Effects of parameter “keep-basic”: Licenses – not deleted; profile – reset; configuration – reset; management IP – reset • Effects of parameter “keep-virt-vols”: Licenses – deleted; profile – reset; configuration – reset; management IP – unchanged • Confirming the command causes system reboot	

reset factory keep-docker

	reset factory keep-docker Resets all switch configuration except docker configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.7.11xx
Role	Admin
Example	`(config) # reset factory keep-docker`
Related Commands	reset factory reset factory halt reset factory keep-all-config reset factory keep-basic reset factory keep-virt-vols reset factory only-config
Notes	N/A

Configuration new factory keep-docker

	Configuration new <config_file_name> factory keep-docker Creates new file with only factory defaults except docker current configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.7.11xx
Role	Admin
Example	`(config) # no configuration new my_file factory keep-docker`
Related Commands	configuration new <config_file_name> factory configuration new <config_file_name> factory keep-basic configuration new <config_file_name> factory keep-connect
Notes	N/A

show files debug-dump

show files debug-dump [<filename>]

Displays a list of debug dump files.

Syntax Description	filename	Displays a summary of the contents of a particular debug dump file.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch [standalone: master] (config) # show files debug-dump sysdump-switch-20170731-161038.tgz`
	`switch [standalone: master] (config) # show files debug-dump sysdump-switch-20170731-161038.tgz`
	`==========`
	`System information:`
	
	`Hostname:  switch`
	`Version:  X86_64 3.6.4006 2017-07-03 16:17:39 x86_64`
	`Current time:  2017-07-31 16:10:38`
	`System uptime: 19d 18h 20m 12s`
	
	`==========`
	
	`==========`
	`Output of 'uname -a':`
	
	`Linux switch 3.10.0-327.36.3.el7smp-x86_64 X86_64 jenkins #1 2017-06-27 12:34:55 SMP x86_64 x86_64 x86_64 GNU/Linux`
	
	`==========`	

Related Commands	file debug-dump	
Notes		

show files stats

show files stats <filename>

Displays a list of statistics report files.

Syntax Description	filename	Display the contents of a particular statistics report file.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show files stats`
	`memory-201140524-111745.csv`
	`switch (config) #`	

Related Commands	file stats	
Notes		

show files system

show files system [detail]

Displays usage information of the file systems on the system.

Syntax Description	detail	Displays more detailed information on file-system
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show files system`
	`Statistics for /config filesystem:`
	`  Bytes Total  100 MB`
	`  Bytes Used  3 MB`
	`  Bytes Free  97 MB`
	`  Bytes Percent Free  97%`
	`  Bytes Available  97 MB`
	`  Inodes Total  0`
	`  Inodes Used  0`
	`  Inodes Free  0`
	`  Inodes Percent Free  0%`
	
	`Statistics for /var filesystem:`
	`  Bytes Total  860 MB`
	`  Bytes Used  209 MB`
	`  Bytes Free  651 MB`
	`  Bytes Percent Free  75%`
	`  Bytes Available  651 MB`
	`  Inodes Total  0`
	`  Inodes Used  0`
	`  Inodes Free  0`
	`  Inodes Percent Free  0%`
	`switch (config) #`	

Related Commands	N/A	
Notes		

show files tcpdump

	show files tcpdump Displays a list of statistics report files.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show files stats`
	`test`
	`dump3`
	`switch (config) #`

Related Commands	file tcpdump tcpdump
Notes	

4.6.5.2 Configuration Files

configuration audit

configuration audit max-changes <number>

Chooses settings related to configuration change auditing.

Syntax Description	max-changes	Set maximum number of audit messages to log per change.
Default	1000	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # configuration audit max-changes 100`
	`switch (config) # show configuration audit`
	`Maximum number of changes to log: 100`
	`switch (config) #`	

Related Commands	show configuration	
Notes	N/A	

configuration copy

configuration copy <source name> <dest name>

Copies a configuration file.

Syntax Description	source name	Name of source file.
	dest name	Name of destination file. If the file of specified file-name does not exist a new file will be created with said filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration copy initial.bak example` `switch (config) #`	
Related Commands		
Notes	• This command does not affect the current running configuration • The active configuration file may not be the target of a copy. However, it may be the source of a copy in which case the original remains active.	

configuration delete

configuration delete <filename>

Deletes a configuration file.

Syntax Description	filename	Name of file to delete.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show configuration files`
	`example  initial  initial.bak  initial.prev`
	`switch (config) # configuration delete example`
	`switch (config) # show configuration files`
	`initial  initial.bak  initial.prev`
	`switch (config) #`	

Related Commands	show configuration	
Notes	• This command does not affect the current running configuration • The active configuration file may not be deleted	

configuration fetch

configuration fetch <URL> [<name>]

Downloads a configuration file from a remote host.

Syntax Description	URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
	name	The configuration file name.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration fetch scp://root:password@ 192.168.10.125/tmp/conf1` `switch (config) #`	
Related Commands	configuration switch-to	
Notes	• The downloaded file should not override the active configuration file, using the <name> parameter • If no name is specified for a configuration fetch, it is given the same name as it had on the server • No configuration file may have the name “active”	

configuration jump-start

	configuration jump-start Runs the initial-configuration wizard.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # configuration jump-start`
	`Mellanox configuration wizard`
	`Step 1: Hostname? [switch-3cc29c]`
	`Step 2: Use DHCP on mgmt0 interface? y`
	`Step 3: Admin password (Enter to leave unchanged)?`
	`You have entered the following information:`
	`1. Hostname: switch-3cc29c`
	`2. Use DHCP on mgmt0 interface: yes`
	`3. Enable IPv6: yes`
	`4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes`
	`53. Admin password (Enter to leave unchanged): (unchanged)`
	`To change an answer, enter the step number to return to.`
	`Otherwise hit <enter> to save changes and exit.`
	`Choice:`
	`Configuration changes saved.`
	`switch (config) #`

Related Commands	N/A
Notes	• The wizard is automatically invoked whenever the CLI is launched when the active configuration file is fresh (i.e. not modified from its initial contents) • This command invokes the wizard on demand – see chapter “Initializing the Switch for the First Time” in the *Onyx User Manual*

configuration merge

configuration merge <filename>

Merges the “shared configuration” from one configuration file into the running configuration.

Syntax Description	filename	Name of file from which to merge settings
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration merge new-config-file` `switch (config) #`	
Related Commands		
Notes	• No configuration files are modified during this process • The configuration filename must be a non-active configuration file	

configuration move

configuration move <source name> <dest name>

Moves a configuration file.

Syntax Description	source name	Old name of file to move.
	dest name	New name for moved file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show configuration files`
	`example1 initial initial.bak initial.prev`
	`switch (config) # configuration move example1 example2`
	`switch (config) # show configuration files`
	`example2 initial initial.bak initial.prev`
	`switch (config) #`	

Related Commands	show configuration	
Notes	• This command does not affect the current running configuration • The active configuration file may not be the target of a move	

configuration new

configuration new <filename> [factory [keep-basic] [keep-connect]]

Creates a new configuration file under the specified name. The parameters specify what configuration, if any, to carry forward from the current running configuration.

Syntax Description	filename	Names for new configuration file.
	factory	Creates new file with only factory defaults.
	keep-basic	Keeps licenses and host keys.
	keep-connect	Keeps configuration necessary for connectivity (interfaces, routes, and ARP).
Default	Keeps licenses and host keys	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show configuration files`
	`initial  initial.bak  initial.prev`
	`switch (config) # configuration new example2`
	`switch (config) # show configuration files`
	`example2  initial  initial.bak  initial.prev`
	`switch (config) #`	

Related Commands	show configuration	
Notes		

configuration switch-to

configuration switch-to <filename> [no-reboot]

Loads the configuration from the specified file and makes it the active configuration file.

Syntax Description	no-reboot	Forces configuration change without rebooting the switch
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.6.1002	Added “no-reboot” option
Role	admin	

Example	`switch (config) # show configuration files`
	`initial (active)`
	`newcon`
	`initial.prev`
	`initial.bak`
	`switch (config) # configuration switch-to newcon no-reboot`
	`switch (config) # show configuration files`
	`initial`
	`newcon (active)`
	`initial.prev`
	`initial.bak`
	`switch (config) #`	

Related Commands	show configuration files	
Notes	• The current running configuration is lost and not automatically saved to the previous active configuration file. • When running the command without the “no-reboot” parameter, the user is prompted to OK a reboot. If the answer is “yes”, the configuration is replaced and the switch is rebooted immediately.	

configuration text fetch

configuration text fetch <URL> [apply [discard | fail-continue | filename | overwrite | verbose] | filename <filename> | overwrite [apply | filename <filename>]]

Fetches a text configuration file (list of CLI commands) from a specified URL.

Syntax Description	apply	Applies the file to the running configuration (i.e. executes the commands in it). This option has the following parameters: • discard: Does not keep downloaded configuration text file after applying it to the system • fail-continue: If applying commands, continues execution even if one of them fails • overwrite: If saving the file and the filename already exists, replaces the old file • verbose: Displays all commands being executed and their output instead of just those that get errors
	filename	Specifies filename for saving downloaded text file.
	overwrite	Downloads the file and saves it using the same name it had on the server. This option has the following parameters: • apply: Applies the downloaded configuration to the running system • filename: Specifies filename for saving downloaded text file
Default	N/A	
Configuration Mode	config	
History	3.2.1000	
	3.2.3000	Updated command
Role	admin	
Example	`switch (config) # configuration fetch text scp://username[:password]@hostname/path/filename`	
Related Commands	N/A	
Notes		

configuration text file

configuration text file <filename> {apply [fail-continue] [verbose] | delete | rename <filename> | upload < URL>}

Performs operations on text-based configuration files.

Syntax Description	filename <file>	Specifies the filename
	apply	Applies the configuration on the system
	fail-continue	Continues execution of the commands even if some commands fail
	verbose	Displays all commands being executed and their output, instead of just those that get errors
	delete	Deletes the file
	rename <filename>	Renames the file
	upload <URL>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration text file my-config-file delete` `switch (config) #`	
Related Commands	show configuration files	
Notes		

configuration text generate

configuration text generate {active {running | saved} | file <filename> } {save <filename> | upload <URL>}

Generates a new text-based configuration file from this system's configuration.

Syntax Description	active	Generates from currently active configuration.
	running	Uses running configuration.
	saved	Uses saved configuration.
	file <filename>	Generates from inactive saved configuration.
	save	Saves new file to local persistent storage.
	upload <URL>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # configuration text generate file initial.prev save example`
	`switch (config) # show configuration files`
	`initial (active)`
	`initial.prev`
	`initial.bak`
	`Active configuration: initial`
	`Unsaved changes:  yes`
	`switch (config) #`	

Related Commands	show configuration files	
Notes		

configuration upload

configuration upload {active | <name>} <URL or scp or sftp://username:password@hostname[:port]/path/filename>

Uploads a configuration file to a remote host.

Syntax Description	active	Upload the active configuration file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration upload active scp://root:password@192.168.10.125/tmp/conf1` `switch (config) #`	
Related Commands	N/A	
Notes	No configuration file may have the name “active”.	

configuration write

configuration write [local | to <filename> [no-switch]]

Saves the running configuration to the active configuration file.

Syntax Description	local	Saves the running configuration locally (same as “write memory local”)
	to <filename>	Saves the running configuration to a new file under a different name and makes it the active file
	no-switch	Saves the running configuration to this file but keep the current one active
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # configuration write` `switch (config) #`	
Related Commands	write	
Notes		

write

write {memory [local] | terminal}

Saves or displays the running configuration.

Syntax Description	memory	Saves running configuration to the active configuration file. It is the same as “configuration write”.
	local	Saves the running configuration only on the local node. It is the same as “configuration write local”.
	terminal	Displays commands to recreate current running configuration. It is the same as “show running-config”.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # write terminal`
	`##`
	`## Running database "initial"`
	`## Generated at 20114/05/27 10:05:16 +0000`
	`## Hostname: switch`
	`##`
	`##`
	`## Network interface configuration`
	`##`
	`interface mgmt0 comment ""`
	`interface mgmt0 create`
	`interface mgmt0 dhcp`
	`interface mgmt0 display`
	`interface mgmt0 duplex auto`
	`interface mgmt0 mtu 1500`
	`no interface mgmt0 shutdown`
	`interface mgmt0 speed auto`
	`no interface mgmt0 zeroconf`
	`##`
	`## Local user account configuration`
	`##`
	`username a** capability admin`
	`no username a** disable`
	`username a** disable password`
	`......`
	`switch (config) #`	

Related Commands	show running-config configuration write	
Notes		

show configuration

show configuration [audit | files [<filename>] | running | text files]

Displays a list of CLI commands that will bring the state of a fresh system up to match the current persistent state of this system.

Syntax Description	audit	Displays settings for configuration change auditing.
	files [<filename>]	Displays a list of configuration files in persistent storage if no filename is specified. If a filename is specified, it displays the commands to recreate the configuration in that file. In the latter case, only non-default commands are shown, as for the normal "show configuration" command.
	running	Displays commands to recreate current running configuration. Same as "show configuration" except that it applies to the currently running configuration, rather than the current persisted configuration.
	text files	Displays names of available text-based configuration files.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.5006	Removed "running full" and "full" parameters
Role	monitor/admin	

Example

```
switch (config) # show configuration
##
## Active saved database "newcon"
## Generated at 20114/05/25 10:18:52 +0000
## Hostname: switch-3cc29c
##
##
## Network interface configuration
##
interface mgmt0 comment ""
interface mgmt0 create
interface mgmt0 dhcp
interface mgmt0 display
interface mgmt0 duplex auto
interface mgmt0 mtu 1500
no interface mgmt0 shutdown
interface mgmt0 speed auto
no interface mgmt0 zeroconf
switch (config) #
```

Related Commands
Notes

show running-config

show running-config [expanded | protocol <protocol>]

Displays commands to recreate current running configuration.

Syntax Description	expanded	Displays commands in expanded format without compressing ranges
	protocol	Only displays commands relating to the specified protocol
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4402	Removed “full” parameter
	3.6.2002	Updated Example and added parameters
	3.6.3640	Added support for forwarding mode configuration
Role	monitor/admin	

Example

```
switch (config) # show running-config
##
## Running database "initial"
## Generated at 2018/07/25 19:34:11 +0000
## Hostname: switch
##

##
## Running-config temporary prefix mode setting
##
no cli default prefix-modes enable

##
## License keys
##
   license install <license>

##
## Other IP configuration
##
   hostname switch

##
## Local user account configuration
##
   username root nopassword

##
## AAA remote server configuration
##
# ldap bind-password ********
# radius-server key ********
# tacacs-server key ********

##
## SNMP configuration
##
   snmp-server user 7YLAyJrC77 v3 capability admin
   snmp-server user 7YLAyJrC77 v3 enable
   snmp-server user 7YLAyJrC77 v3 enable sets
no snmp-server user 7YLAyJrC77 v3 require-privacy
   snmp-server user kRg5dmdogX v3 capability admin
   snmp-server user kRg5dmdogX v3 enable
   snmp-server user kRg5dmdogX v3 enable sets
no snmp-server user kRg5dmdogX v3 require-privacy

##
## Network management configuration
##
# web proxy auth basic password ********

##
## Persistent prefix mode setting
##
cli default prefix-modes enable
```

Related Commands

Notes

4.7 Logging

4.7.1 Monitor

➢ ***To print logging events to the terminal:***

Set the modules or events you wish to print to the terminal. For example, run:

```
switch (config) # logging monitor events notice
switch (config) # logging monitor sx-sdk warning
```

These commands print system events in severity “notice” and “sx-sdk” module notifications in severity “warning” to the screen. For example, in case of interface-down event, the following gets printed to the screen.

```
switch (config) #
Wed Jul 10 11:30:42 2013: Interface IB1/17 changed state to DOWN
Wed Jul 10 11:30:43 2013: Interface IB1/18 changed state to DOWN
```

To see a list of the events, refer to Table 26, “Supported Event Notifications and MIB Mapping,” on page 380.

4.7.2 Remote Logging

➢ ***To configure remote syslog to send syslog messages to a remote syslog server:***

Step 1. Enter Config mode. Run:

```
switch >
switch > enable
switch # configure terminal
```

Step 2. Set remote syslog server. Run

```
switch (config) # logging <IP address/hostname>
```

Step 3. (Optional) Set the destination port of the remote host. Run:

```
switch (config) # logging <IP address/hostname> port <port>
```

Step 4. Set the minimum severity of the log level to `info`. Run:

```
switch (config) # logging <IP address/hostname> trap info
```

Step 5. Override the log levels on a per-class basis. Run:

```
switch (config) # logging <IP address/hostname> trap override class <class name> prior-
ity <level>
```

4.7.3 Commands

logging port

logging <syslog IPv4 address/hostname> port <destination-port>
no logging <syslog IPv4 address/hostname> port

Configures remote server destination port for log messages.
The no form of the command resets the remote log port to its default value.

Syntax Description	destination-port	Range: 1-65535
	Hostname	Max 64 characters
Default	514 (UDP)	
Configuration Mode	config	
History	3.6.2002	
Role	admin	
Example	`switch (config) # logging 10.0.0.1 port 105`	
Related Commands	logging <syslog IPv4 address/hostname> trap	
Notes		

logging trap

logging <syslog IPv4 address/hostname> [trap {<log-level> | override class <class> priority <log-level>}]
no logging <syslog IPv4 address/hostname> [trap {<log-level> | override class <class> priority <log-level>}]

Enables (by setting the IPv4 address/hostname) sending logging messages, with ability to filter the logging messages according to their classes.
The no form of the command stops sending messages to the remote syslog server.

Syntax Description	syslog IPv4 address/hostname	IPv4 address/hostname of the remote syslog server. Hostname is limited to 64 characters
	log-level	• alert – alert notification, action must be taken immediately • crit – critical condition • debug – debug level messages • emerg – system is unusable (emergency) • err – error condition • info – informational condition • none – disables the logging locally and remotely • notice – normal, but significant condition • warning – warning condition
	class	Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with "logging local <log level>". Classes that do have an override will do as the override specifies. If "none" is specified for the log level, the software will not log anything from this class. Classes available: • iss-modules – protocol stack • mgmt-back – system management back-end • mgmt-core – system management core • mgmt-front – system management front-end • mlx-daemons – management daemons • sx-sdk – switch SDK
	log-level	• alert – alert notification, action must be taken immediately • crit – critical condition • debug – debug level messages • emerg – system is unusable (emergency) • err – error condition • info – informational condition • none – disables the logging locally and remotely • notice – normal, but significant condition • warning – warning condition

Default	Remote logging is disabled
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # logging local info`
Related Commands	show logging logging local override logging <syslog IPv4 address/hostname> port
Notes	

logging debug-files

logging debug-files {delete {current | oldest} | rotation {criteria | force | max-num} | update {<number> | current} | upload <log-file> <upload URL>}

Configures settings for debug log files.

Syntax Description	delete {current | oldest}	Deletes certain debug-log files. • current: Deletes the current active debug-log file • oldest: Deletes some of the oldest debug-log files

	rotation {criteria {frequency {daily | weekly | monthly} | size <size> | size-pct <percentage>} | force | max-num}	Configures automatic rotation of debug-logging files.
		• criteria: Sets how the system decides when to rotate debug files.
		  • frequency: Rotate log files on a fixed time-based schedule
		  • size: Rotate log files when they pass a size threshold in megabytes
		  • size-pct: Rotate logs when they surpass a specified percentage of disk
		• forces: Forces an immediate rotation of the log files
		• max-num: Specifies the maximum number of old log files to keep

	update {<number> | current}	Uploads a local debug-log file to a remote host. • current: Uploads log file “messages” to a remote host • number: Uploads compressed log file “debug.<number>.gz” to a remote host. Range is 1-10
	upload	Uploads debug log file to a remote host
	log-file	Possible values: 1-7, or current
	upload URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported (e.g.: scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # logging debug-files delete current` `switch (config) #`	
Related Commands		
Notes		

logging event enable

	logging events {cpu-rate-limiters | interfaces | protocols} enable **no logging events {cpu-rate-limiters | interfaces | protocols} enable** Activate event tracking for a certain group. The no form of the command deactivates event tracking for a certain group.
Syntax Description	cpu-rate-limiters | interfaces | protocols — Logical groups with specified set of counters
Default	N/A
Configuration Mode	config
History	3.6.6000
Role	admin
Example	`switch (config) # logging events interfaces enable`
Related Commands	
Notes	

logging event error-threshold

	logging events {cpu-rate-limiters | interfaces | protocols} error-threshold <events> **no logging events {cpu-rate-limiters | interfaces | protocols} error-threshold <events>** Configures number of events after which the system begins to generate events to the log file. The no form of the command resets this parameter to its default value.	
Syntax Description	cpu-rate-limiters	Sets threshold for CPU rate limiter related events
	interfaces	Sets threshold for interface related events
	protocols	Sets threshold for protocol related events
	events	Number of events after which the system begins to generate events to the log file. Range: 0-4294967295.
Default	cpu-rate-limiters – 1 event interfaces – 10 events protocols – 2 events	
Configuration Mode	config	
History	3.6.6000	
Role	admin	
Example	`switch (config) # logging events interfaces error-threshold 45`	
Related Commands		
Notes		

logging event interval

logging events {cpu-rate-limiters | interfaces | protocols} interval <seconds>
no logging events {cpu-rate-limiters | interfaces | protocols} interval <seconds>

Configures interval in seconds between each sampling of counters in event type. The no form of the command resets this parameter to its default value.

Syntax Description	cpu-rate-limiters | interfaces | protocols	Logical groups with specified set of counters
	seconds	Time between sampling. Range is different for each event type: • cpu-rate-limiters – 5-3600 • interfaces – 10-3600 • protocols – 10-3600
Default	cpu-rate-limiters – 10 seconds interfaces – 5 minutes protocols – 1 minute	
Configuration Mode	config	
History	3.6.6000	
Role	admin	
Example	`switch (config) # logging events interfaces interval 120`	
Related Commands		
Notes		

logging event rate-limit

logging events [cpu-rate-limiters | interfaces | protocols] rate-limit {short | medium | long} [count | window]
no logging events [cpu-rate-limiters | interfaces | protocols] rate-limit [short | medium | long] [count <number> | window <seconds>]

Configures the number of allowed events per time window and that window's duration.
The no form of the command resets these parameters to their default values.

Syntax Description

cpu-rate-limiters | interfaces | protocols	Logical groups with specified set of counters
rate-limit	Three configurable periods: short, medium, and long
count	Number of allowed events per time window
window	Window of time in seconds for the rate limit period

Default

For "interfaces"	For "protocols"	For "cpu-rate-limiters"
Short window: event count – 5 window duration – 1 hour	Short window: event count – 10 window duration – 1 hour	Short window: event count – 10 window duration – 1 hour
Medium window: event count – 50 window duration – 1 day	Medium window: event count – 100 window duration – 1 day	Medium window: event count – 200 window duration – 1 day
Long window: event count – 350 window duration – 7 days	Long window: event count – 600 window duration – 7 days	Long window: event count – 1200 window duration – 7 days

Configuration Mode

config

History

3.6.6000

Role

admin

Example

```
switch (config) # logging events interfaces interval 120
```

Related Commands

Notes

- The goal of this command is to restrict the number of events in the log. To achieve this end, it is possible to specify the allowed number (parameter "count") of messages per period of time (parameter "window").

logging fields

logging fields seconds {enable | fractional-digits <f-digit> | whole-digits <w-digit>}
no logging fields seconds {enable | fractional-digits <f-digit> | whole-digits <w-digit>}

Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not.
The no form of the command disallows including an additional field in each log message that shows the number of seconds since the Epoch.

Syntax Description	enable	Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not.
	f-digit	The fractional-digits parameter controls the number of digits to the right of the decimal point. Truncation is done from the right. Possible values are: 1, 2, 3, or 6.
	w-digit	The whole-digits parameter controls the number of digits to the left of the decimal point. Truncation is done from the left. Except for the year, all of these digits are redundant with syslog's own date and time. Possible values: 1, 6, or all.
Default	disabled	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # logging fields seconds enable
switch (config) # logging fields seconds whole-digits 1
switch (config) # show logging
Local logging level: info
  Override for class mgmt-front: warning
Default remote logging level: notice
No remote syslog servers configured.
Allow receiving of messages from remote hosts: no
Number of archived log files to keep: 10
Log rotation size threshold: 5.000% of partition (43 megabytes)
Log format: standard
Subsecond timestamp field: enabled
Subsecond timestamp precision: 1 whole digit; 3 fractional digits
Levels at which messages are logged:
  CLI commands: notice
  Audit messages: notice
switch (config) #
```

Related Commands	show logging
Notes	This is independent of the standard syslog date and time at the beginning of each message in the format of “July 15 18:00:00”. Aside from indicating the year at full precision, its main purpose is to provide subsecond precision.

logging files delete

logging files delete {current | oldest [<number of files>]}

Deletes the current or oldest log files.

Syntax Description	current	Deletes current log file.
	oldest	Deletes oldest log file.
	number of files	Sets the number of files to be deleted.
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # logging files delete current` `switch (config) #`	
Related Commands	show logging show log files	
Notes		

logging files rotation

logging files rotation {criteria { frequency <freq> | size <size-mb>| size-pct <size-percentage>} | force | max-number <number-of-files>}

Sets the rotation criteria of the logging files.

Syntax Description	freq	Sets rotation criteria according to time. Possible options are: • Daily • Weekly • Monthly
	size-mb	Sets rotation criteria according to size in mega bytes. The range is 1-9999.
	size-percentage	Sets rotation criteria according to size in percentage of the partition where the logging files are kept in. The percentage given is truncated to three decimal points (thousandths of a percent).
	force	Forces an immediate rotation of the log files. This does not affect the schedule of auto-rotation if it was done based on time: the next automatic rotation will still occur at the same time for which it was previously scheduled. Naturally, if the auto-rotation was based on size, this will delay it somewhat as it reduces the size of the active log file to zero.
	number-of-files	The number of log files will be kept. If the number of log files ever exceeds this number (either at rotation time, or when this setting is lowered), the system will delete as many files as necessary to bring it down to this number, starting with the oldest.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # logging files rotation criteria size-pct 6`
	`switch (config) # show logging`
	`Local logging level: info`
	`  Override for class mgmt-front: warning`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: no`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 6.000% of partition (51.60 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: enabled`
	`Subsecond timestamp precision: 1 whole digit; 3 fractional digits`
	`Levels at which messages are logged:`
	`  CLI commands: info`
	`  Audit messages: notice`
	`switch (config)`

Related Commands	show logging show log files
Notes	

logging files upload

logging files upload {current | <file-number>} <url>

Uploads a log file to a remote host.

Syntax Description	current	The current log file. The current log file will have the name “messages” if you do not specify a new name for it in the upload URL.
	file-number	An archived log file. The archived log file will have the name “messages<n>.gz” (while “n” is the file number) if you do not specify a new name for it in the upload URL. The file will be compressed with gzip.
	url	Uploads URL path. FTP, TFTP, SCP, and SFTP are supported. For example: scp://username[:password]@hostname/path/filename.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # logging files upload 1 scp://admin@scpserver`	
Related Commands	show logging show log files	
Notes		

logging format

logging format {standard | welf [fw-name <hostname>]}
no logging format {standard | welf [fw-name <hostname>]}

Sets the format of the logging messages.
The no form of the command resets the format to its default.

Syntax Description	standard	Standard format.
	welf	WebTrends Enhanced Log file (WELF) format.
	hostname	Specifies the firewall hostname that should be associated with each message logged in WELF format. If no firewall name is set, the hostname is used by default. hostname is limited to 64 characters.
Default	standard	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # logging format standard`
	`switch (config) # show logging`
	`Local logging level: info`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: yes`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 5.000% of partition (43 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: disabled`
	`Levels at which messages are logged:`
	`  CLI commands: notice`
	`  Audit messages: notice`
	`switch (config) #`	

Related Commands	show logging	
Notes		

logging level

logging level {cli commands <log-level> | audit mgmt <log-level>}

Sets the severity level at which CLI commands or the management audit message that the user executes are logged. This includes auditing of both configuration changes and actions.

Syntax Description	cli commands	Sets the severity level at which CLI commands which the user executes are logged.
	audit mgmt	Sets the severity level at which all network management audit messages are logged.
	log-level	• alert – alert notification, action must be taken immediately • crit – critical condition • debug – debug level messages • emerg – system is unusable (emergency) • err – error condition • info – informational condition • none – disables the logging locally and remotely • notice – normal, but significant condition • warning – warning condition
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # logging level cli commands info`
	`switch (config) # show logging`
	`Local logging level: info`
	`  Override for class mgmt-front: warning`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: no`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 5.000% of partition (43 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: enabled`
	`Subsecond timestamp precision: 1 whole digit; 3 fractional digits`
	`Levels at which messages are logged:`
	`  CLI commands: info`
	`  Audit messages: notice`
	`switch (config) #`	

Related Commands	show logging	
Notes		

logging local override

logging local override [class <class> priority <log-level>]
no logging local override [class <class> priority <log-level>]

Enables class-specific overrides to the local log level.
The no form of the command disables all class-specific overrides to the local log level without deleting them from the configuration, but disables them so that the logging level for all classes is determined solely by the global setting.

Syntax Description	override	Enables class-specific overrides to the local log level.
	class	Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with “logging local <log level>”. Classes that do have an override will do as the override specifies. If “none” is specified for the log level, the software will not log anything from this class. Classes available: • debug-module - debug module functionality • protocol-stack - protocol stack modules functionality • mgmt-back - system management back-end components • mgmt-core - system management core • mgmt-front - system management front-end components • mlx-daemons - management daemons • sx-sdk - switch SDK
	log-level	• alert - alert notification, action must be taken immediately • crit - critical condition • debug - debug level messages • emerg - system is unusable (emergency) • err - error condition • info - informational condition • none - disables the logging locally and remotely • notice - normal, but significant condition • warning - warning condition
Default	Override is disabled.	
Configuration Mode	config	
History	3.1.0000	
	3.3.4150	Added debug-module class Changed iss-modules with protocol-stack
Role	admin	

Example	`switch (config) # logging local override class mgmt-front priority warning`
	`switch (config) # show logging`
	`Local logging level: info`
	`  Override for class mgmt-front: warning`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: no`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 5.000% of partition (43 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: disabled`
	`Levels at which messages are logged:`
	`  CLI commands: notice`
	`  Audit messages: notice`
	`switch (config) #`

Related Commands	show logging logging local
Notes	

logging monitor

logging monitor <facility> <priority-level>
no logging monitor <facility> <priority-level>

Sets monitor log facility and level to print to the terminal.
The no form of the command disables printing logs of facilities to the terminal.

Syntax Description	facility	• mgmt-front • mgmt-back • mgmt-core • events • sx-sdk • mlnx-daemons • iss-modules
	priority-level	• none • emerg • alert • crit • err • warming • notice • info • debug
Default	no logging monitor	
Configuration Mode	config	
History	3.3.4000	
Role	admin	
Example	`switch (config) # logging monitor events notice` `switch (config) #`	
Related Commands		
Notes		

logging receive

	logging receive **no logging receive** Enables receiving logging messages from a remote host. The no form of the command disables the option of receiving logging messages from a remote host.
Syntax Description	N/A
Default	Receiving logging is disabled
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # logging receive`
	`switch (config) # show logging`
	`Local logging level: info`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: yes`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 5.000% of partition (43 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: disabled`
	`Levels at which messages are logged:`
	`  CLI commands: notice`
	`  Audit messages: notice`
	`switch (config) #`

Related Commands	show logging logging local logging local override
Notes	• This does not log to the console TTY port • In-band management should be enabled in order to open a channel from the host to the CPU • If enabled, only log messages matching or exceeding the minimum severity specified with the “logging local” command will be logged, regardless of what is sent from the remote host

logging trap

logging trap <log-level>
no logging trap

Configures the minimum severity of log messages sent to syslog servers.
The no form of the command disables sending event log messages to syslog servers.

Syntax Description	log-level	The minimum severity level for all configured syslog servers: • none – disable logging • emerg – emergency: system is unusable • alert – action must be taken immediately • crit – critical conditions • err – error conditions • warning – warning conditions • notice – normal but significant condition • info – informational messages • debug – debug-level messages
Default	Receiving logging is disabled	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # logging trap info` `switch (config) #`	
Related Commands		
Notes		

show log

show log [continuous | files [<file-number>]] [[not] matching <reg-exp>]

Displays the log file with optional filter criteria.

Syntax Description	continues	Displays the last few lines of the current log file and then continues to display new lines as they come in until the user hits Ctrl+C, similar to LINUX “tail” utility.
	files	Displays the list of log files.
	<file-number>	Displays an archived log file, where the number may range from 1 up to the number of archived log files available.
	[not] matching <reg-exp>	The file is piped through a LINUX “grep” utility to only include lines either matching, or not matching, the provided regular expression.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.3.4402	Updated example and added note
Role	admin	

Example	`switch [standalone: master] (config) # show log matching "Executing|Action"`
	`Jul 31 16:11:23 M2100-aj cli[26502]: [cli.NOTICE]: user : Executing command: enable`
	`Jul 31 16:11:24 M2100-aj cli[26507]: [cli.NOTICE]: user : Executing command: enable`
	`Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: enable`
	`Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: show license`
	`Jul 31 16:11:41 M2100-aj cli[26548]: [cli.NOTICE]: user : Executing command: enable`
	`Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: enable`
	`Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: conf termina`	

Related Commands	logging fields logging files rotation logging level logging local logging receive show logging	
Notes	• When using a regular expression containing | (OR), the expression should be surrounded by quotes (“<expression>”), otherwise it is parsed as filter (PIPE) command. • The command’s output has many of the options as the Linux “less” command. These options allow navigating the log file and perform searches. To see help for different option press “h” after running the “show log” command.	

show logging

	show logging Displays the logging configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show logging`
	`Local logging level: info`
	`  Override for class mgmt-front: warning`
	`Default remote logging level: notice`
	`No remote syslog servers configured.`
	`Allow receiving of messages from remote hosts: no`
	`Number of archived log files to keep: 10`
	`Log rotation size threshold: 5.000% of partition (43 megabytes)`
	`Log format: standard`
	`Subsecond timestamp field: enabled`
	`Subsecond timestamp precision: 1 whole digit; 3 fractional digits`
	`Levels at which messages are logged:`
	`  CLI commands: info`
	`  Audit messages: notice`
	`switch (config) #`

Related Commands	logging fields logging files rotation logging level logging local logging receive logging <syslog IPv4 address/hostname>
Notes	

show logging events

show logging events [cpu-rate-limiters | interfaces | protocols]

Displays configuration per selected event group or all.

Syntax Description	cpu-rate-limiters | interfaces | protocols — Logical groups with specified set of counters
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000
Role	admin
Example	

```
switch (config) # show logging events

cpu-rate-limiters:
  Admin mode     : yes
  Interval       : 10 seconds
  Error threshold: 1

  Rate-limit short window:
    Event count    : 10
    Window duration: 1 hour

  Rate-limit medium window:
    Event count    : 200
    Window duration: 1 day

  Rate-limit long window:
    Event count    : 1200
    Window duration: 7 days

interfaces:
  Admin mode     : no
  Interval       : 5 minutes
  Error threshold: 10

  Rate-limit short window:
    Event count    : 5
    Window duration: 1 hour

  Rate-limit medium window:
    Event count    : 50
    Window duration: 1 day

  Rate-limit long window:
    Event count    : 350
    Window duration: 7 days
```

```
protocols:
  Admin mode     : no
  Interval       : 1 minute
  Error threshold: 2

  Rate-limit short window:
    Event count    : 10
    Window duration: 1 hour

  Rate-limit medium window:
    Event count    : 100
    Window duration: 1 day

  Rate-limit long window:
    Event count    : 600
    Window duration: 7 days
```

Related Commands

Notes

show logging events source-counters

	show logging events [cpu-rate-limiters | interfaces | protocols] source-counters Displays set of counters for sampling.
Syntax Description	cpu-rate-limiters | interfaces | protocols — Logical groups with specified set of counters
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000
Role	admin

Example	`switch (config) # show logging events interfaces source-counters`
	
	`interfaces:`
	`  Counters: Rx discard packets, Rx error packets, Rx fcs errors, Rx undersize packets, Rx oversize packets, Rx unknown control opcode, Rx symbol errors, Rx discard packets by Storm Control, Tx discard packets, Tx error packets, Tx hoq discard packets`

Related Commands	
Notes	

4.8 Debugging

➢ ***To use the debugging logs feature:***

Step 1. Enable debugging. Run:

```
switch (config) # debug ethernet all
```

Step 2. Display the debug level set. Run:

```
switch (config) # show debug ethernet
```

Step 3. Display the logs. Run:

```
switch (config) # show log debug {match | continue}
```

4.8.1 Commands

debug ethernet all

debug ethernet all
no debug ethernet all

Enables debug traces for Ethernet modules.
The no form of the command disables the debug traces for all Ethernet modules.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4150
Role	admin
Example	`switch (config) # debug ethernet all` `switch (config) #`
Related Commands	
Notes	

debug ethernet dcbx

debug ethernet dcbx {all | management | fail-all | control-panel | tlv}

Configures the trace level for DCBX.
The no form of the command disables the configured DCBX debug traces.

Syntax Description	all	Enables all traces.
	management	Management messages.
	fail-all	All failure traces.
	control-panel	Control plane traces.
	tlv	TLV related trace configuration.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet dcbx all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet ip igmp-snooping

debug ethernet ip igmp-snooping {all | forward-db-messages | group-info | init-shut | packet-dump | query | source-info | system-resources-management | timer | vlan-info}
no debug ethernet ip igmp-snooping {all | forward-db-messages | group-info | init-shut | packet-dump | query | source-info | system-resources-management | timer | vlan-info}

Configures the trace level for IGMP snooping.
The no form of the command disables tracking a specified level.

Syntax Description	all	Enable track traces
	forward-db-messages	Forwarding database messages
	group-info	Group information messages
	init-shut	Init and shutdown messages
	packet-dump	Packet dump messages
	query	Query related messages
	source-info	Source information messages
	system-resources-management	System resources management messages
	timer	Timer messages
	vlan-info	VLAN information messages
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet ip igmp-snooping all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet ip interface

debug ethernet ip interface {all | arp-packet-dump | buffer | enet-packet-dump | error | fail-all | filter | trace-error | trace-event}
no debug ethernet ip interface {all | arp-packet-dump | buffer | enet-packet-dump | error | fail-all | filter | trace-error | trace-event}

Configures the trace level for interface.
The no form of the command disables tracking a specified level.

Syntax Description	all	Enable track traces
	arp-packet-dump	ARP packet dump trace
	buffer	Buffer trace
	enet-packet-dump	ENET packet dump trace
	error	Trace error messages
	fail-all	All failures including Packet Validation Trace
	filter	Lower layer traces
	trace-error	Trace error messages
	trace-event	Trace event messages
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet ip interface all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet lacp

debug ethernet lacp {all | all-resource | data-path | fail-all | init-shut | management | memory | packet}
no debug ethernet lacp {all | all-resources | data-path | fail-all | init-shut | management | memory | packet}

Configures the trace level for LACP.
The no form of the command disables the configured LACP debug traces.

Syntax Description	all	Enables all traces.
	all-resource	BPDU related messages.
	data-path	Init and shutdown traces.
	fail-all	Management messages.
	init-shut	Memory related messages.
	management memory	IP packet dump trace.
	memory	All failure traces.
	packet	OS resource trace.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet lacp all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet lldp

debug ethernet lldp {all | control-panel | critical-event | data-path | fail-all | init-shut | management | memory | neigh-add | neigh-age-out | neigh-del | neigh-drop | neigh-updt | tlv}
no debug ethernet lldp {all | control-panel | critical-event | data-path | fail-all | init-shut | management | memory | neigh-add | neigh-age-out | neigh-del | neigh-drop | neigh-updt | tlv}

Configures the trace level for LLDP.
The no form of the command disables the configured LLDP debug traces.

Syntax Description	all	Enables all traces.
	control-panel	Control plane traces.
	critical-event	Critical traces.
	data-path	IP packet dump trace.
	fail-all	All failure traces.
	init-shut	Init and shutdown traces.
	management	Management messages.
	memory	Memory related messages.
	neigh-add	Neighbor add traces.
	neigh-age-out	Neighbor ageout traces.
	neigh-del	Neighbor delete traces.
	neigh-drop	Neighbor drop traces.
	neigh-updt	Neighbor update traces.
	tlv	TLV related trace configuration
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet lldp all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet port

	debug ethernet port all Configures the trace level for port. The no form of the command disables the configured port debug traces.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4150
Role	admin
Example	`switch (config) # debug ethernet port all` `switch (config) #`
Related Commands	
Notes	

debug ethernet qos

debug ethernet qos {all | all-resource | control-panel | fail-all | filters | init-shut | management | memory | packet}
no debug ethernet qos {all | all-resource | control-panel | fail-all | filters | init-shut | management | memory | packet}

Configures the trace level for QoS.
The no form of the command disables the configured QoS debug traces.

Syntax Description	all	Enables all traces.
	all-resource	OS resource traces.
	control-panel	Control plane traces.
	fail-all	All failure traces.
	filters	Lower layer traces.
	init-shut	Init and shutdown traces.
	management	Management messages.
	memory	Memory related messages.
	packet	BPDU related messages.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet port all` `switch (config) #`	
Related Commands		
Notes		

debug ethernet spanning-tree

debug ethernet spanning-tree {all | error | event | filters | init-shut | management | memory | packet | port-info-state-machine | port-receive-state-machine | port-role-selection-state-machine | port-transit-state-machine | port-transmit-state-machine | protocol-migration-state-machine | timers}
no debug ethernet spanning-tree {all | error | event | filters | init-shut | management | memory | packet | port-info-state-machine | port-receive-state-machine | port-role-selection-state-machine | port-transit-state-machine | port-transmit-state-machine | protocol-migration-state-machine | timers}

Configures the trace level for spanning-tree.
The no form of the command disables the configured spanning-tree debug traces.

Syntax Description	all	Enables all traces.
	error	Error messages trace.
	event	Events related messages.
	filters	Lower later traces.
	init-shut	Init and shutdown traces.
	management	Management messages.
	memory	Memory related messages.
	packet	BPDU related messages.
	port-info-state-machine	Port information messages.
	port-receive-state-machine	Port received messages.
	port-role-selection-state-machine	Port role selection messages.
	port-transit-state-machine	Port transition messages.
	port-transmit-state-machine	Port transmission messages.
	protocol-migration-state-machine	Protocol migration messages.
	timers	Timer modules message.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	

Example	`switch (config) # debug ethernet spanning-tree all` `switch (config) #`
Related Commands	
Notes	

debug ethernet vlan

debug ethernet vlan {all | fwd | priority | filters}
no debug ethernet vlan {all | fwd | priority | filters}

Configures the trace level for VLAN.
The no form of the command disables the configured VLAN debug traces.

Syntax Description	all	Enables all traces
	fwd	Forward.
	priority	Priority.
	filters	Lower layer traces.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config) # debug ethernet vlan all` `switch (config) #`	
Related Commands		
Notes		

show debug ethernet

show debug ethernet {dcbx | ip {arp | dhcp-relay | igmp-snooping | interface | ospf} | lacp | lldp | port | qos | spanning-tree | vlan}

Displays debug level configuration on a specific switch.

Syntax Description	dcbx	Displays the trace level for spanning tree
	ip	Displays debug trace level for ethernet routing module. • arp • dhcp-relay • igmp-snooping • interface • ospf
	lacp	Displays the trace level for LACP
	lldp	Displays the trace level for LLDP
	port	Displays the trace level for port
	qos	Displays the trace level for QoS
	spanning-tree	Displays the trace level for spanning tree
	vlan	Displays the trace level for VLAN
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.6000	Updated Example
Role	admin	

Example	`switch (config) # show debug ethernet dcbx`
	`dcbx protocol:`
	`  management   : ON`
	`  fail-all     : ON`
	`  control-panel: ON`
	`  tlv          : ON`	

Related Commands		
Notes		

show log debug

show log debug [continuous | files | matching | not]

Displays current event debug-log file in a scrollable pager.

Syntax Description	continuous	Displays new event log messages as they arrive.
	files	Displays archived debug log files.
	matching	Displays event debug logs that match a given regular expression.
	not	Displays event debug logs that do not meet certain criteria.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
Role	admin	
Example		
Related Commands		
Notes		

```
switch (config) # show log debug
Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:47 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>QoSHwQueueDelete i4IfIndex[137]
Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:47 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>QoSHwQueueDelete i4IfIndex[141]
Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:48 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: ==FsHwSetSpeed sx_api_port_speed_admin_set = 0
Jun 15 16:20:48 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: ==FsHwGetSpeed sx_api_port_speed_oper_get = 0
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[89], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[33], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[73], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[121], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[133], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[13], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[81], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[117], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[65], u1ConfigOption[6]
.
.
.
switch (config) #
```

4.9 Link Diagnostic Per Port

4.9.1 General

When debugging a system, it is important to be able to quickly identify the root of a problem. The Diagnostic commands enables an insight into the physical layer components where the user is able to see information such as a cable status (plugged/unplugged) or if Auto-Negotiation has failed.

4.9.2 List of Possible Output Messages

```
No issue was observed
Closed by command
Negotiation failure
Link training failure
Speed logical mismatch
Remote faults detected
Cable speed not enabled
Bad signal integrity
Other issues
Speed degradation
Information unavailable
Cable is unplugged
Unsupported cable
I2C bus is stuck
Module memory invalid
Module overheated
Module short circuit
Power budget exceeded
Management forced down
```

Commands

show interfaces ethernet link-diagnostics

show interfaces ethernet [<interface>] link-diagnostics

Displays a specific Ethernet module/port or all Ethernet ports.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4006
	3.6.4110 Updated command input.
Role	admin
Example	`switch (config) # show interfaces ethernet link-diagnostics`

```
--------------------------------------------------------------------------
 Interface   Code        Status
--------------------------------------------------------------------------
 Eth1/1      1024        Cable is unplugged
 Eth1/2      1024        Cable is unplugged
 Eth1/3      1024        Cable is unplugged
 Eth1/4      1024        Cable is unplugged
 Eth1/5      1024        Cable is unplugged
 Eth1/6      1024        Cable is unplugged
 Eth1/7      1024        Cable is unplugged
 Eth1/8      1024        Cable is unplugged
 Eth1/9      1024        Cable is unplugged
 Eth1/10     1024        Cable is unplugged
 Eth1/11     1024        Cable is unplugged
 Eth1/12     1024        Cable is unplugged
 Eth1/13     1024        Cable is unplugged
 Eth1/14     1024        Cable is unplugged
 Eth1/15     1024        Cable is unplugged
 Eth1/16     1024        Cable is unplugged
 Eth1/17     1024        Cable is unplugged
 Eth1/18     1024        Cable is unplugged
 Eth1/19     1024        Cable is unplugged
 Eth1/20     1024        Cable is unplugged
 Eth1/21     1024        Cable is unplugged
 Eth1/22     1024        Cable is unplugged
 Eth1/23     1024        Cable is unplugged
 Eth1/24     1024        Cable is unplugged
 Eth1/25     1024        Cable is unplugged
 Eth1/26     1024        Cable is unplugged
 Eth1/27     1024        Cable is unplugged
 Eth1/28     1024        Cable is unplugged
 Eth1/29     1024        Cable is unplugged
 Eth1/30     1024        Cable is unplugged
 Eth1/31     0           No issue was observed
 Eth1/32     0           No issue was observed
```

Related Commands
Notes

4.10 Signal Degradation Monitoring

A system can monitor the Bit Error Rate (BER) in order to ensure a quality of the link. As long as BER observed by the MAC layer is low enough, the rate of packet loss is low enough to allow successful operation of the applications running on top of the network.

The system continuously monitors the link BER and compares it to BER limits, when limits are crossed the system can generate an event indicating that link quality is degraded to the network operator that can take preemptive actions or even disable the low quality link.

When Forward Error Correction (FEC) is enabled a network operator can choose to monitor an amount of corrected errors by using the pre-FEC mode, or the amount of errors which the FEC failed to correct (uncorrectable errors) by using the post-FEC mode, when FEC is used then every error detected by the PHY will be monitored.

When link is disabled the system will keep it in shutdown state until the port is explicitly enabled (Explicitly running "shutdown" and then "no shutdown" commands for that port).

4.10.1 Effective-BER Monitoring

Effective-BER is the BER that the MAC/Application layer observe. Errors monitored by the Effective-BER may directly result in a packet drop. For links with no error correction, the Effective BER is the BER received by port, and it is monitored based on the received Phy symbols. For links with FEC, the Effective BER represents the rate of errors that the FEC decoder did not manage to correct and were passed to the MAC layer. The Effective BER for FEC links is monitored using the FEC decoder uncorrectable codewords data.

4.10.2 Configuring Signal Degradation Monitoring

Step 1. Enable signal degradation monitoring. Run:

```
switch (config) #  1/3 signal-degrade
```

If not indicated, the interface is disabled in case of signal degradation.

Step 2. (Optional) To prevent the interface from shutting down in case of signal degradation, run:

```
switch (config) #  1/3 signal-degrade no-shutdown
```

Step 3. (Optional) Enable SNMP notifications on signal degradation events. Run:

```
switch (config) # snmp notify event health-module-status
```

Please refer to Section 4.18.1.7, “Configuring an SNMP Notification,” on page 571 for a general explanation on how to enable SNMP notifications for specific events.

Step 4. (Optional) Enable email notifications on signal degradation events. Run:

```
switch (config) # email notify event health-module-status
```

Please refer to Section 4.11.3, “Email Notifications,” on page 383 for a general explanation on how to enable email notifications for specific events.

4.10.3 Commands

signal-degrade

<slot>/<port> signal-degrade [no-shutdown]
no <slot>/<port> signal-degrade [no-shutdown]

Enables signal degradation operation per interface.

Syntax Description	no-shutdown	Does not shutdown an affected interface
Default	Disabled	
Configuration Mode	config	
History	3.6.4110	
Role	admin	
Example	`switch (config) # 1/1 signal-degrade`	
Related Commands		
Notes		

show interfaces ethernet signal-degrade

show interfaces ethernet [<slot>/<port>] signal-degrade

Displays signal degradation information.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4110
Role	admin

Example	`switch (config) # show interfaces ethernet signal-degrade`
	`---------------------------------------------------------------------------------------`
	`Interface  Admin state                  Monitoring   Action       FEC type`
	`---------------------------------------------------------------------------------------`
	`Eth1/1      Enabled                      Disabled     Shutdown     no-fec/post-fec`
	`Eth1/2      Enabled                      Disabled     Shutdown     no-fec/post-fec`
	`Eth1/3      Enabled                      Disabled     Shutdown     no-fec/post-fec`
	`Eth1/4      Enabled                      Disabled     Shutdown     no-fec/post-fec`
	`Eth1/5      Enabled                      Disabled     Shutdown     no-fec/post-fec`
	`...`

Related Commands	
Notes	

4.11 Event Notifications

Onyx features a variety of supported events. Events are printed in the system log file and can, optionally, be sent to the system administrator via email, SNMP trap or directly prompted to the terminal.

4.11.1 Supported Events

Table 26 presents the supported events and maps them to their relevant MIB OID.

Table 26 - Supported Event Notifications and MIB Mapping

Event Name	Event Description	MIB OID	Comments
asic-chip-down	ASIC (chip) down	Mellanox-EFM-MIB: asicChipDown	Not supported
cpu-util-high	CPU utilization has risen too high	Mellanox-EFM-MIB: cpuUtilHigh	N/A
dcbx-ets-port-admin-state-trap	DCBX ETS port admin state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortAdminStateTrap	N/A
dcbx-ets-port-oper-state-trap	DCBX ETS port oper state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortOperStateTrap	N/A
dcbx-ets-port-peer-state-trap	DCBX ETS port peer state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortPeerStateTrap	N/A
dcbx-pfc-module-state-change	DCBX PFC module state change	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCModuleStateTrap	N/A
dcbx-pfc-port-admin-state-trap	DCBX PFC port admin state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortAdminStateTrap	N/A
dcbx-pfc-port-oper-state-trap	DCBX PFC port oper state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortOperStateTrap	N/A
dcbx-pfc-port-peer-state-trap	DCBX PFC port peer state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortPeerStateTrap	N/A
disk-space-low	File system free space has fallen too low	Mellanox-EFM-MIB: diskSpaceLow	N/A
health-module-status	Health module status changed	Mellanox-EFM-MIB: systemHealthStatus	N/A
insufficient-fans	Insufficient amount of fans in system	Mellanox-EFM-MIB: insufficientFans	N/A
insufficient-fans-recover	Insufficient amount of fans in system recovered	Mellanox-EFM-MIB: insufficientFansRecover	N/A
insufficient-power	Insufficient power supply	Mellanox-EFM-MIB: insufficientPower	N/A

Table 26 - Supported Event Notifications and MIB Mapping

Event Name	Event Description	MIB OID	Comments
interface-down	An interface's link state has changed to DOWN	RFC1213: linkdown (SNMPv1)	Supported for Ethernetand management interfaces for 1U and blade systems
interface-up	An interface's link state has changed to UP	RFC1213: linkup (SNMPv1)	Supported for Ethernetand management interfaces for 1U and blade systems
internal-bus-error	Internal bus (I^2C) error	Mellanox-EFM-MIB: internalBusError	N/A
liveness-failure	A process in the system is detected as hung	Not implemented	N/A
low-power	Low power supply	Mellanox-EFM-MIB: lowPower	N/A
low-power-recover	Low power supply recover	Mellanox-EFM-MIB: lowPowerRecover	N/A
mstp-new-bridge-root	The bridge become the root bridge root of a MSTI	MELLANOX-MSTP-MIB: mstpRootBridgeChange	N/A
mstp-new-root-port	The root port of a MSTI changed	MELLANOX-MSTP-MIB: mstpRootPortChange	N/A
mstp-topology-change	Port in MSTI become forwarding of blocking	MELLANOX-MSTP-MIB: mstpTopologyChange	N/A
N/A	Reset occurred due to over-heating of ASIC	Mellanox-EFM-MIB: asicOverTempReset	Not supported
new_root	Local bridge became a root bridge	Bridge-MIB: newRoot	N/A
ospf-auth-fail	OSPF authentication failure	OSPF-TRAP-MIB: ospfIfAuthFailure	N/A
ospf-config-error	OSPF config error	OSPF-TRAP-MIB: ospfIfConfigError	N/A
ospf-if-rx-bad-packet	Bad OSPF packet received	OSPF-TRAP-MIB: ospfIfRxBadPacket	N/A
ospf-if-state-change	OSPF interface state change	OSPF-TRAP-MIB: ospfIfStateChange	N/A
ospf-lsdb-approaching-overflow	OSPF LSDB is approaching overflow	OSPF-TRAP-MIB: ospfLsdbApproachingOverflow	Not supported
ospf-lsdb-overflow	OSPF LSDB overflow	OSPF-TRAP-MIB: ospfLsdbOverflow	Not supported

Table 26 - Supported Event Notifications and MIB Mapping

Event Name	Event Description	MIB OID	Comments
ospf-nbr-state-change	OSPF neighbor state change	OSPF-TRAP-MIB: ospfNbrStateChange	N/A
paging-high	Paging activity has risen too high	N/A	Not supported
process-crash	A process in the system has crashed	Mellanox-EFM-MIB: procCrash	N/A
process-exit	A process in the system unexpectedly exited	Mellanox-EFM-MIB: procUnexpectedExit	N/A
send-test	Send a test notification	testTrap	Run CLI command # snmp-server notify send-test
snmp-authtrap	An SNMPv3 request has failed authentica-tion	Not implemented	N/A
temperature-too-high	Temperature is too high	Mellanox-EFM-MIB: asicOverTemp	N/A
topology_change	Topology change trig-gered by a local bridge	Bridge-MIB: topologyChange	N/A
unexpected-shutdown	Unexpected system shutdown	Mellanox-EFM-MIB: unexpectedShutdown	N/A
xstp-new-root-bridge	The bridge became the root bridge of STI	MELLANOX-XSTP-MIB: mellanoxXstpRootBridgeChange	N/A
xstp-root-port-change	XSTP root port changed	MELLANOX-XSTP-MIB: mellanoxXstpRootPortChange	N/A
xstp-topology-change	Port in pvrst become forwarding of blocking	MELLANOX-XSTP-MIB: mellanoxXstpTopologyChange	N/A

4.11.2 Terminal Notifications

➢ ***To print events to the terminal:***

Set the events you wish to print to the terminal. Run:

```
switch (config) # logging monitor events notice
```

This command prints system events in the severity “notice” to the screen. For example, in case of interface-down event, the following gets printed to the screen.

```
switch (config) #
Wed Jul 10 11:30:42 2013: Interface IB1/17 changed state to DOWN
Wed Jul 10 11:30:43 2013: Interface IB1/18 changed state to DOWN
switch (config) #
```

4.11.3 Email Notifications

➢ ***To configure Onyx to send you emails for all configured events and failures:***

Step 1. Enter to Config mode. Run:

```
switch >
switch > enable
switch # configure terminal
```

Step 2. Set your mailhub to the IP address to be your mail client's server – for example, Microsoft Outlook exchange server.

```
switch (config) # email mailhub <IP address>
```

Step 3. Add your email address for notifications. Run:

```
switch (config) # email notify recipient <email address>
```

Step 4. Configure the system to send notifications for a specific event. Run:

```
switch (config) # email notify event <event name>
```

Step 5. Show the list of events for which an email is sent. Run:

```
switch (config) # show email events
Failure events for which emails will be sent:
  process-crash: A process in the system has crashed
  unexpected-shutdown: Unexpected system shutdown

Informational events for which emails will be sent:
  asic-chip-down: ASIC (Chip) Down
  cpu-util-high: CPU utilization has risen too high
  cpu-util-ok: CPU utilization has fallen back to normal levels
  disk-io-high: Disk I/O per second has risen too high
  disk-io-ok: Disk I/O per second has fallen back to acceptable levels
  disk-space-low: Filesystem free space has fallen too low
.
.
.
switch (config) #
```

Step 6. Have the system send you a test email. Run:

```
switch # email send-test

The last command should generate the following email:
-----Original Message-----
From: Admin User [mailto:do-not-reply@switch.]
Sent: Sunday, May 01, 2011 11:17 AM
To: <name>
Subject: System event on switch: Test email for event notification

==== System information:
Hostname: switch
Version:  <version> 2011-05-01 14:56:31
```

```
          ...
Date:     2011/05/01 08:17:29
Uptime:   17h 8m 28.060s

This is a test email.
==== Done.
```

4.11.4 Commands

4.11.4.1 Email Notification

email autosupport enable

email autosupport enable
no email autosupport enable

Sends automatic support notifications via email.
The no form of the command stops sending automatic support notifications via email.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # email autosupport enable`
Related Commands	N/A
Notes	

email autosupport event

email autosupport event <event>
no email autosupport event

Specifies for which events to send auto-support notification emails.
The no form of the command resets auto-support email security mode to its default.

Syntax Description	event	• process-crash – a process has crashed • process-exit – a process unexpectedly exited • liveness-failure – a process iss detected as hung • cpu-util-high – CPU utilization has risen too high • cpu-util-ok – CPU utilization has fallen back to normal levels • paging-high – paging activity has risen too high • paging-ok – paging activity has fallen back to normal levels • disk-space-low – filesystem free space has fallen too low • disk-space-ok – filesystem free space is back in the normal range • memusage-high – memory usage has risen too high • memusage-ok – memory usage has fallen back to acceptable levels • netusage-high – network utilization has risen too high • netusage-ok – network utilization has fallen back to acceptable levels • disk-io-high – disk I/O per second has risen too high • disk-io-ok – disk I/O per second has fallen back to acceptable • levels • unexpected-cluster-join – node has unexpectedly joined the cluster • unexpected-cluster-leave – node has unexpectedly left the cluster • unexpected-cluster-size – the number of nodes in the cluster is unexpected • unexpected-shutdown – unexpected system shutdown • interface-up – an interface's link state has changed to up • interface-down – an interface's link state has changed to down • user-login – a user has logged into the system • user-logout – a user has logged out of the system • health-module-status – health module status • temperature-too-high – temperature has risen too high • low-power – low power supply • low-power-recover – low power supply recover • insufficient-power – insufficient power supply • power-redundancy-mismatch – power redundancy mismatch • insufficient-fans – insufficient amount of fans in system • insufficient-fans-recover – insufficient amount of fans in system recovered

	• asic-chip-down – ASIC (chip) down • internal-bus-error – internal bus (I^2C) error • internal-link-speed-mismatch – internal links speed mismatch
Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # email autosupport event process-crash`
Related Commands	N/A
Notes	

email autosupport ssl mode

email autosupport ssl mode {none | tls | tls-none}
no email autosupport ssl mode

Configures type of security to use for auto-support email.
The no form of the command resets auto-support email security mode to its default.

Syntax Description	none	Does not use TLS to secure auto-support email.
	tls	Uses TLS over the default server port to secure auto-support email and does not send an email if TLS fails.
	tls-none	Attempts TLS over the default server port to secure auto-support email, and falls back on plaintext if this fails.
Default	tls-none	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # email autosupport ssl mode tls`	
Related Commands	N/A	
Notes		

email autosupport ssl cert-verify

	email autosupport ssl cert-verify **no email autosupport ssl cert-verify** Verifies server certificates. The no form of the command does not verify server certificates.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # email autosupport ssl cert-verify`
Related Commands	N/A
Notes	

email autosupport ssl ca-list

email autosupport ssl ca-list {<ca-list-name> | default_ca_list | none}
no email autosupport ssl ca-list

Configures supplemental CA certificates for verification of server certificates. The no form of the command removes supplemental CA certificate list.

Syntax Description	default_ca_list	Default supplemental CA certificate list.
	none	No supplemental list; uses built-in list only.
Default	default_ca_list	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # email autosupport ssl ca-list default_ca_list`	
Related Commands	N/A	
Notes		

email dead-letter

email dead-letter {cleanup max-age <duration> | enable}
no email dead-letter

Configures settings for saving undeliverable emails.
The no form of the command disables sending of emails to vendor auto-support upon certain failures.

Syntax Description	duration	Example: “5d4h3m2s” for 5 days, 4 hours, 3 minutes, 2 seconds.
	enable	Saves dead-letter files for undeliverable emails.
Default	Save dead letter is enabled The default duration is 14 days	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # email dead-letter enable` `switch (config) #`	
Related Commands	show email	
Notes		

email domain

email domain <hostname or IP address>
no email domain

Sets the domain name from which the emails will appear to come from (provided that the return address is not already fully-qualified). This is used in conjunction with the system hostname to form the full name of the host from which the email appears to come.
The no form of the command clears email domain override.

Syntax Description	hostname or IP address IP address.
Default	No email domain
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # email domain mellanox`
	`switch (config) # show email`
	`Mail hub: 10.0.8.11`
	`Mail hub port: 125`
	`Domain: mellanox`
	`Return address: do-not-reply`
	`Include hostname in return address: yes`
	`...`
	`switch (config) #`

Related Commands	show emails
Notes	

email mailhub

email mailhub <hostname or IP address>
no email mailhub

Sets the mail relay to be used to send notification emails.
The no form of the command clears the mail relay to be used to send notification emails.

Syntax Description	hostname or IP address	Hostname or IP address.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # email mailhub 10.0.8.11`
	`switch (config) # show email`
	`Mail hub: 10.0.8.11`
	`Mail hub port: 25`
	`Domain: (not specified)`
	`Return address: do-not-reply`
	`Include hostname in return address: yes`
	`...`
	`switch (config) #`	

Related Commands	show email [events]	
Notes		

email autosupport mailhub

email autosupport mailhub <hostname or IP address>
no email autosupport mailhub

Sets the mail relay to be used for sending autosupport notification emails. The no form of the command clears the mail relay to be used for sending autosupport notification emails.

Syntax Description	<Hostname or IP address> The mail hub Hostname or IP address.
Default	N/A
Configuration Mode	config
History	3.7.1000
Role	Admin

Example	`switch (config) # email autosupport mailhub 10.10.10.1`
	`switch (config) # show email`
	
	`Autosupport emails`
	`  Enabled:  no`
	`  Recipient:`
	`  Mail hub:  10.10.10.1`
	`  Security mode:  tls-none`
	`  Verify server cert:  yes`
	`  Supplemental CA list: default-ca-list`

Related Commands	show email
Notes	

email autosupport recipient

email autosupport recipient <email addr>
no email autosupport recipient

Sets the recipient for autosupport emails.
The no form of the command clears the configured autosupport recipient.

Syntax Description	<email addr>	The autosupport recipient email address.
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	Admin	

Example	`switch (config) # email autosupport recipient user@example.com`
	`switch (config) # show email`
	
	`Autosupport emails`
	`  Enabled:  no`
	`  Recipient:  user@example.com`
	`  Mail hub:`
	`  Security mode:  tls-none`
	`  Verify server cert:  yes`
	`  Supplemental CA list: default-ca-list`	

Related Commands	show email	
Notes		

email mailhub-port

email mailhub-port <hostname or IP address>
no email mailhub-port

Sets the mail relay port to be used to send notification emails.
The no form of the command resets the port to its default.

Syntax Description	hostname or IP address hostname or IP address.
Default	25
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # email mailhub-port 125`
	`switch (config) # show email`
	`Mail hub: 10.0.8.11`
	`Mail hub port: 125`
	`Domain: (system domain name)`
	`Return address: do-not-reply`
	`Include hostname in return address: yes`
	`...`
	`switch (config) #`

Related Commands	show email
Notes	

email notify event

email notify event <event name>
no email notify event <event name>

Enables sending email notifications for the specified event type.
The no form of the command disables sending email notifications for the specified event type.

Syntax Description	event name	Example event names would include “process-crash” and “cpu-util-high”.
Default	No events are enabled	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # email notify event process-crash`
	`switch (config) # show email events`
	`Failure events for which emails will be sent:`
	`process-crash: A process in the system has crashed`
	`unexpected-shutdown: Unexpected system shutdown`
	
	`Informational events for which emails will be sent:`
	`liveness-failure: A process in the system was detected as hung`
	`process-exit: A process in the system unexpectedly exited`
	`cpu-util-ok: CPU utilization has fallen back to normal levels`
	`cpu-util-high: CPU utilization has risen too high`
	`disk-io-ok: Disk I/O per second has fallen back to acceptable levels`
	`...`
	`temperature-too-high: Temperature has risen too high`
	
	`All events for which autosupport emails will be sent:`
	`process-crash: A process in the system has crashed`
	`liveness-failure: A process in the system was detected as hungswitch`
	`(config) #`
	`switch (config) #`	

Related Commands	show email	
Notes	This does not affect auto-support emails. Auto-support can be disabled overall, but if it is enabled, all auto-support events are sent as emails.	

email notify recipient

email notify recipient <email addr> [class {info | failure} | detail]
no email notify recipient <email addr> [class {info | failure} | detail]

Adds an email address from the list of addresses to which to send email notifications of events.
The no form of the command removes an email address from the list of addresses to which to send email notifications of events.

Syntax Description	email addr	Email address of intended recipient.
	class	Specifies which types of events are sent to this recipient.
	info	Sends informational events to this recipient.
	failure	Sends failure events to this recipient.
	detail	Sends detailed event emails to this recipient.
Default	No recipients are added	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # email notify recipient user2@autosupport.mellanox.com`
	`switch (config) # show email`
	`Mail hub:`
	`Mail hub port: 25`
	`Domain: (not specified)`
	`Return address: user1`
	`Include hostname in return address: no`
	`Dead letter settings:`
	`Save dead.letter files: yes`
	`Dead letter max age: (none)`
	`Email notification recipients:`
	`user2@autosupport.mellanox.com (all events, in detail)`
	`Autosupport emails`
	`Enabled: no`
	`Recipient: autosupport@autosupport.mellanox.com`
	`Mail hub: autosupport.mellanox.com`
	`switch (config) #`	

Related Commands	show email	
Notes		

email return-addr

email return-addr <username>
no email domain

Sets the username or fully-qualified return address from which email notifications are sent.

- If the string provided contains an "@" character, it is considered to be fully-qualified and used as-is.
- Otherwise, it is considered to be just the username, and we append "@<hostname>.<domain>". The default is "do-not-reply", but this can be changed to "admin" or whatnot in case something along the line does not like fictitious addresses.

The no form of the command resets this attribute to its default.

Syntax Description	username	Username.
Default	do-not-reply	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # email return-addr user1`
	`switch (config) # show email`
	`Mail hub:`
	`Mail hub port: 25`
	`Domain: (not specified)`
	`Return address: user1`
	`Include hostname in return address: yes`
	`...`
	`switch (config) #`	

Related Commands	show email	
Notes		

email return-host

	email return-host **no email return-host** Includes the hostname in the return address for emails. The no form of the command does not include the hostname in the return address for emails.
Syntax Description	N/A
Default	No return host
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # no email return-host`
	`switch (config) # show email`
	`Mail hub:`
	`Mail hub port:    25`
	`Domain:           (system domain name)`
	`Return address:   my-address`
	`Include hostname in return address: no`
	
	`Current reply address: host@localdomain`
	
	`Dead letter settings:`
	`  Save dead.letter files: yes`
	`  Dead letter max age:    5 days`
	
	`No recipients configured.`
	
	`Autosupport emails`
	`  Enabled:      no`
	`  Recipient:    autosupport@autosupport.mellanox.com`
	`  Mail hub:     autosupport.mellanox.com`
	`switch (config) #`

Related Commands	show email
Notes	This only takes effect if the return address does not contain an “@” character.

email send-test

	email send-test Sends test-email to all configured event and failure recipients.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # email send-test`
Related Commands	show email [events]
Notes	

email ssl mode

email ssl mode {none | tls | tls-none}
no email ssl mode

Sets the security mode(s) to try for sending email.
The no form of the command resets the email SSL mode to its default.

Syntax Description	none	No security mode, operates in plaintext.
	tls	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it gives up.
	tls-none	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it falls back on plaintext.
Default	default-cert	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # email ssl mode tls-none`	
Related Commands	N/A	
Notes		

email ssl cert-verify

	email ssl cert-verify **no email ssl cert-verify** Enables verification of SSL/TLS server certificates for email. The no form of the command disables verification of SSL/TLS server certificates for email.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # email ssl cert-verify`
Related Commands	N/A
Notes	This command has no impact unless TLS is used.

email ssl ca-list

email ssl ca-list {<ca-list-name> | default-ca-list | none}
no email ssl ca-list

Specifies the list of supplemental certificates of authority (CA) from the certificate configuration database that is to be used for verification of server certificates when sending email using TLS, if any.
The no form of the command uses no list of supplemental certificates.

Syntax Description	ca-list-name	Specifies CA list name.
	default-ca-list	Uses default supplemental CA certificate list.
	none	Uses no list of supplemental certificates.
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # email ssl ca-list none`	
Related Commands	N/A	
Notes	This command has no impact unless TLS is used, and certificate verification is enabled.	

show email

show email [events]

Displays email configuration or events for which email should be sent upon.

Syntax Description	events	show event list
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Role	admin	

Example	`switch (config) # show email`
	
	`Autosupport emails`
	`  Enabled:  no`
	`  Recipient:`
	`  Mail hub:  10.10.10.1`
	`  Security mode:  tls-none`
	`  Verify server cert:  yes`
	`  Supplemental CA list: default-ca-list`	

Related Commands	show email	
Notes		

4.12 Telemetry

As it is becoming increasingly complex to manage networks, and network administrators need more tools to understand network behavior, it is necessary to provide basic information about network performance, identify network bottlenecks, and provide information for the purposes of network optimization and future planning.

Therefore, network administrators are required to constantly review network port behavior, record port buffer consumption, and identify shortage in buffer resources and record flows which lead to the excessive buffer consumption.

Onyx provides the following mechanisms to perform these tasks:

- Sampling (histograms) – a network administrator can enable a sampling of the port buffer occupancy, record occupancy changes over time, and provide information for different levels of buffer occupancy, and amount of time the buffer has been occupied during the observation period.
- Thresholds – thresholds may be enabled per port to record the network time when port buffer occupancy crosses the defined threshold and when buffer occupancy drops below it.
- Flow recording – a record of the most active flows which cause an excessive usage of the port buffers may be kept. Once enabled, the system may identify flow patterns and present a user with a list of flows, based on which a network administrator can rearrange distribution of the data flows in the network and minimize data loss.

4.12.1 Commands

protocol telemetry

protocol telemetry
no protocol telemetry

Unhides telemetry config CLIs.
The no form of the command hides telemetry config CLIs.

Syntax Description	N/A
Default	Disabled.
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # protocol telemetry` `switch (config) # no protocol telemetry`
Related Commands	
Notes	

telemetry shutdown

	telemetry shutdown **no telemetry shutdown** Disables the telemetry protocol, threshold detection, and histogram fetching for all sampling enabled interfaces without changing any internal configuration. The no form of the command enables telemetry protocol.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # telemetry shutdown` `switch (config) # no telemetry shutdown`
Related Commands	
Notes	

telemetry sampling log

telemetry sampling log <time>
no telemetry sampling log <time>

Enables the log interval value (histogram fetching) from device.
The no form of the command disables the log interval value.

Syntax Description	time	Input Range: 100 msec - 1 min
Default	1000 msec.	
Configuration Mode	config	
History	3.6.3004	
Role	admin	
Example	`switch (config) # telemetry sampling log 1000` `switch (config) # no telemetry sampling log`	
Related Commands		
Notes		

telemetry sampling tc

telemetry sampling tc <0-7> [mcast | ucast]
no telemetry sampling tc <tc_id> [mcast | ucast]

Enables multicast sampling (histogram fetching) on a traffic class for a particular Ethernet interface.
The no form of the command disables multicast sampling on a TC for a particular Ethernet interface.

Syntax Description	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.3004	
Role	admin	
Example	`switch (config  1/2) # telemetry sampling tc 3 mcast`	
Related Commands		
Notes		

telemetry threshold

telemetry threshold tc <0-7> [ucast | mcast]
no telemetry threshold tc <0-7> [ucast | mcast]

Enables threshold in hardware for a particular traffic class.
The no form of the command disables threshold in hardware for a particular traffic class.

Syntax Description	ucast	Unicast traffic
	mcast	Multicast traffic
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/12) # telemetry threshold tc 0 ucast`	
Related Commands		
Notes		

telemetry threshold level

telemetry threshold level <level>
no telemetry threshold level <level>

Configures the threshold level in the hardware per port.
The no form of the command resets the parameter to its default.

Syntax Description	level	Range: 96-1,000,000 Level is set in bytes and in increments of 96
Default	69984	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/12) # telemetry threshold level 288`	
Related Commands		
Notes		

telemetry threshold log

	telemetry threshold log **no telemetry threshold log** Enables logging of threshold events in syslog. The no form of the command disable logging.
Syntax Description	N/A
Default	false
Configuration Mode	config
History	3.6.4006
Role	admin
Example	`switch (config) # telemetry threshold log` `switch (config) # no telemetry threshold log`
Related Commands	
Notes	

telemetry threshold syslog

telemetry threshold syslog <time>
no telemetry threshold syslog <time>

The command sets threshold events logging rate on per hour basis.
The no form of the command sets the logging rate back to default.

Syntax Description	time	Max rate per hour. Input range: 1-3600
Default	100	
Configuration Mode	config	
History	3.6.4006	
Role	admin	
Example	`switch (config) # telemetry threshold syslog 400`	
Related Commands		
Notes		

clear telemetry

clear telemetry {threshold | sampling} [interface <type> <port-id>] [tc <0-7> [ucast | mcast]]

Clears telemetry data.

Syntax Description	type	Possible values: ethernet, port-channel, mlag-port-channel
	tc	Traffic class
	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/12) # clear telemetry threshold level 288`	
Related Commands		
Notes		

clear telemetry threshold

clear telemetry threshold [interface <type> <if>]

Clears threshold and top talker data.

Syntax Description	type	Possible values: ethernet, port-channel, mlag-port-channel
Default	N/A	
Configuration Mode	config	
History	3.6.6105	
Role	admin	
Example	`switch (config) # clear telemetry threshold interface ethernet 1/34-1/36`	
Related Commands		
Notes		

stats export csv telemetry

stats export csv telemetry <slot>/<port>[/<subport>]/<tc>-[mcast | ucast] [file-name *] [after * *] [before * *]

Exports histograms collected by stats to a csv file.

Syntax Description	slot/port	Port number
	subport	Sub-port number to be used in case of split port
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # stats export csv telemetry 1/1` `Generated report file: telemetry-20170119-102715.csv`	
Related Commands		
Notes		

file stats telemetry delete

	file stats telemetry delete <filename> Deletes the given .csv file created by “stats export” command to user directory.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # file stats telemetry delete telemetry-20171006-102158.csv`
Related Commands	
Notes	

file stats telemetry upload

file stats telemetry upload <filename> <upload-url>

Uploads .csv file created by “stats export” command to user directory.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # file stats telemetry upload telemetry-20170119-102715.csv scp://username:password@server//directory` `Password (if required): ******`
Related Commands	
Notes	

show telemetry

show telemetry

Displays the global configuration of telemetry properties.

Syntax Description	N/A
Default	
Configuration Mode	config
History	3.6.4000
Role	admin

Example

```
switch (config) # show telemetry
Telemetry Status                  : Enabled
H/W Sampling Interval(nsec)       : 512
S/W Sampling Interval(ms)         : 1000
Threshold Logging                 : Disabled
Threshold Logging(rate per hour)  : 100

-------------------------------------------------------------------------------------------
Interface         TC                Sampling          Threshold         Level (bytes)
-------------------------------------------------------------------------------------------
Eth1/1            0 (ucast)         Enabled           Disabled          N/A
Eth1/1            5 (mcast)         Enabled           Disabled          N/A
Eth1/2            0 (ucast)         Enabled           Disabled          N/A
Eth1/2            5 (mcast)         Enabled           Disabled          N/A
Eth1/3            N/A               Disabled          Disabled          200 (192)
Eth1/4            3 (mcast)         Disabled          Enabled           69984 (69984)
Eth1/5            3 (mcast)         Disabled          Enabled           69984 (69984)
Eth1/6            N/A               Disabled          Disabled          N/A
Eth1/7            N/A               Disabled          Disabled          N/A
...
Eth1/32           Disabled          Disabled          Disabled          N/A
```

Related Commands	
Notes	

show telemetry sampling tc mcast

show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> mcast

Displays fetched multicast histogram details for a given tc_id of the Ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show telemetry sampling 1/2 tc 3 mcast`	

```
Telemetry histogram: Eth1/2 traffic-class 3 - mcast
       Time                                Bin sizes (nsec buffer was occupied in bytes range)
-------------------------------------------------------------------------------------------------------------------------
01/16/17            2976<           27552     52128     76704     101280    125856    150432    175008    199584    199584>
04:09:07.79936  1000000000          0         0         0         0         0         0         0         0         0
04:09:08.80096  1000000000          0         0         0         0         0         0         0         0         0
04:09:09.80355  1000000000          0         0         0         0         0         0         0         0         0
04:09:10.80518  1000000000          0         0         0         0         0         0         0         0         0
04:09:11.80682  1000000000          0         0         0         0         0         0         0         0         0
```

Related Commands	
Notes	

show telemetry sampling tc mcast last

show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> mcast last <num_of_entries>

Displays last num of fetched multicast histogram details for the given tc_id of the ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
	num_of_entries	Input range: 0-1000
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show telemetry sampling 1/2 tc 3 mcast last 4`	

```
Telemetry histogram: Eth1/2 traffic-class 3 - mcast
       Time                                  Bin sizes (nsec buffer was occupied in bytes range)
-------------------------------------------------------------------------------------------------------------------------
01/16/17              2976<          27552     52128     76704     101280    125856    150432    175008    199584    199584>
04:23:38.28864   1000000000          0         0         0         0         0         0         0         0         0
04:23:39.28977   1000000000          0         0         0         0         0         0         0         0         0
04:23:40.29111   1000000000          0         0         0         0         0         0         0         0         0
04:23:41.29259   1000000000          0         0         0         0         0         0         0         0         0
```

Related Commands	
Notes	In case requested entries are more than what the DB contains it will print the amount in the table.

show telemetry sampling tc ucast

show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> ucast

Displays fetched unicast histogram details for a given TC ID of the Ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show telemetry sampling 1/2 tc 6 ucast`	
Related Commands		
Notes		

```
Telemetry histogram: Eth1/2 traffic-class 6 - ucast
       Time                                Bin sizes (nsec buffer was occupied in bytes range)
-------------------------------------------------------------------------------------------------------------------
01/13/17            2976<           27552     52128     76704     101280    125856    150432    175008    199584    199584>
08:18:09.67745   1000000000         0         0         0         0         0         0         0         0         0
08:18:10.67850   1000000000         0         0         0         0         0         0         0         0         0
08:18:11.67953   1000000000         0         0         0         0         0         0         0         0         0
```

show telemetry sampling tc ucast last

show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> ucast last <num_of_entries>

Displays last num of fetched unicast histogram details for the given tc_id of the ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
	num_of_entries	Input range: 0-1000
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show telemetry sampling 1/2 tc 3 ucast last 3`	

```
Telemetry histogram: Eth1/2 traffic-class 3 - ucast
       Time                                  Bin sizes (nsec buffer was occupied in bytes range)
-------------------------------------------------------------------------------------------------------------------------
01/16/17              2976<           27552     52128     76704     101280    125856    150432    175008    199584    199584>
04:28:39.81351   1000000000           0         0         0         0         0         0         0         0         0
04:28:40.81512   1000000000           0         0         0         0         0         0         0         0         0
04:28:41.81708   1000000000           0         0         0         0         0         0         0         0         0
```

Related Commands	
Notes	In case requested entries are more than what the DB contains it will print the amount in the table.

show telemetry threshold

show telemetry threshold [interface <type> <port-id>] [tc <0-7> [ucast | mcast]]

Displays threshold data for either all interfaces or single interface or per interface per traffic class.

Syntax Description	type	• ethernet • port-channel • mlag-port-channel
	tc	Traffic class
	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6105	Updated Example
Role	admin	

Example

```
switch (config) # show telemetry threshold  1/10-1/13
-------------------------------------------------------------------------------------------------
Event-id   Date      Time        Port        TC     Level       Duration(100 usec)      Repeated
-------------------------------------------------------------------------------------------------
1          09/21/17  10:11:48    Eth 1/10    0      100         102497.61               1
2          09/21/17  10:12:06    Eth 1/10    3      100         85714.76                1
```

```
switch (config) # show telemetry threshold interface port-channel 20 tc 2 mcast
-------------------------------------------------------------------------------------------------------------
Event-id   Date      Time       Port             TC           Level    Duration(100 usec)    Repeated
-------------------------------------------------------------------------------------------------------------
1          09/21/17  10:11:48   Po20 (Eth 1/1)   2 (mcast)    100      102497.61              1
2          09/21/17  10:12:06   Po20 (Eth 1/1)   2 (mcast)    100      85714.76               1
```

Related Commands

Notes

The command supports displaying up to 1000 threshold events. As a result, if more than 1000 thresholds configured in total, some interfaces may not be displayed. Therefore, to query thresholds for a specific interface, please use “show telemetry threshold interface <type> <id>”.

show files stats telemetry

show files stats telemetry [filename]

Displays all files created by the command “stats export csv telemetry”.

Syntax Description	filename	Displays stats for the specified file
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show files stats telemetry telemetry-20180527-102715.csv`
	`Hostname :test-switch`
	`Report :telemetry histogram`
	`Time lower bound(UTC) :2018/05/28 05:58:10`
	`Time upper bound(UTC) :2018/05/28 05:58:25`
	`Export time(UTC) :2018/05/28 06:00:06`
	`Time lower bound :2018/05/28 08:58:10 +0300`
	`Time upper bound :2018/05/28 08:58:25 +0300`
	`Export time :2018/05/28 09:00:06 +0300`
	`System version :X86_64 sys_test 2018-05-15 04:02:13 x86_64`	

Related Commands	stats export csv telemetry	
Notes		

4.13 User Management and Security

4.13.1 User Accounts

There are two general user account types: *admin* and *monitor*. As *admin*, the user is privileged to execute all the available operations. As *monitor*, the user can execute operations that display system configuration and status, or set terminal settings.

Table 27 - User Roles (Accounts) and Default Passwords

User Role	Default Password
admin	admin
monitor	monitor
xmladmin	xmladmin
xmluser	xmluser

To remove passwords from the XML users, run the command `username <username> nopassword`.

4.13.2 Authentication, Authorization and Accounting (AAA)

AAA is a term describing a framework for intelligently controlling access to computer resources, enforcing policies, auditing usage, and providing the information necessary to bill for services. These combined processes are considered important for effective network management and security. The AAA feature allows you to verify the identity of, grant access to, and track the actions of users managing the system. The Onyx switch supports Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) or Lightweight Directory Access Protocol (LDAP) protocols.

- **Authentication** – authentication provides the initial method of identifying each individual user, typically by entering a valid username and password before access is granted. The AAA server compares a user's authentication credentials with the user credentials stored in a database. If the credentials match, the user is granted access to the network or devices. If the credentials do not match, authentication fails and network access is denied.
- **Authorization** – following the authentication, a user must gain authorization for performing certain tasks. After logging into a system, for instance, the user may try to issue commands. The authorization process determines whether the user has the authority to issue such commands. Simply put, authorization is the process of enforcing policies: determining what types or qualities of activities, resources, or services a user is permitted. Usually, authorization occurs within the context of authentication. Once you have authenticated a user, they may be authorized for different types of access or activity.
- **Accounting** – the last level is accounting, which measures the resources a user consumes during access. This includes the amount of system time or the amount of data a user has sent and/or received during a session. Accounting is carried out by logging of session statistics and usage information, and is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Authentication, authorization, and accounting services are often provided by a dedicated AAA server, a program that performs these functions. Network access servers interface with AAA servers using the Remote Authentication Dial-In User Service (RADIUS) protocol.

4.13.2.1 User Re-authentication

Re-authentication prevents users from accessing resources or perform tasks for which they do not have authorization. If credential information (e.g. AAA server information like IP address, key, port number etc.) that has been previously used to authenticate a user is modified, that user gets immediately logged out of the switch and asked to re-authenticate.

4.13.2.2 RADIUS

RADIUS (Remote Authentication Dial-In User Service), widely used in network environments, is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for embedded network devices such as routers, modem servers, switches and so on. RADIUS is currently the de-facto standard for remote authentication. It is prevalent in both new and legacy systems.

It is used for several reasons:

- RADIUS facilitates centralized user administration
- RADIUS consistently provides some level of protection against an active attacker

4.13.2.3 TACACS+

TACACS (Terminal Access Controller Access Control System), widely used in network environments, is a client/server protocol that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for providing NAS (Network Access Security). NAS ensures secure access from remotely connected users. TACACS implements the TACACS Client and provides the AAA (Authentication, Authorization and Accounting) functionalities.

TACACS is used for several reasons:

- Facilitates centralized user administration
- Uses TCP for transport to ensure reliable delivery
- Supports inbound authentication, outbound authentication and change password request for the authentication service
- Provides some level of protection against an active attacker

4.13.2.4 LDAP

LDAP (Lightweight Directory Access Protocol) is an authentication protocol that allows a remote access server to forward a user's log-on password to an authentication server to determine whether access can be allowed to a given system. LDAP is based on a client/server model. The switch acts as a client to the LDAP server. A remote user (the remote administrator) interacts only with the switch, not the back-end server and database.

LDAP authentication consists of the following components:

- A protocol with a frame format that utilizes TCP over IP
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

Each entry in the LDAP server is referenced by its Distinguished Name (DN). The DN consists of the user-account name concatenated with the LDAP domain name. If the user-account name is John, the following is an example DN:

```
uid=John,ou=people,dc=domain,dc=com
```

4.13.3 System Secure Mode

System secure mode is a state that configures the switch system to run secure algorithms in compliance with FIPS 140-2 requirements. In this mode, unsecure algorithms are disabled and unsecure feature configurations are disallowed.

In this mode the system supports Federal Information Processing Standards (FIPS) 140-2, Security Requirements for Cryptographic Modules, which is a NIST (National Institute of Standards and Technology) publication that specifies the requirement for system cypher functionality.

When this mode is activated, all the modules which are used by the system are verified to work in compliance with the secure mode.

Note that if system fails to load in secure mode it is loaded in non-secure mode.

Prerequisites:

Step 1. Disable SNMPv1 and v2. Run:

```
switch (config) # no snmp-server enable communities
```

Step 2. Only allow SNMPv3 users with sha and aes-128. Run:

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128
<password2>
```

Step 3. Only allow SNMPv3 traps with sha and aes-128. Run:

```
switch (config) # snmp-server host <ip-address> informs version 3 user <username> auth
sha <password1> priv aes-128 <password2>
```

Step 4. Only allow SSHv2. Run:

```
switch (config) # ssh server min-version 2
```

Step 5. Enable SSH server strict security mode. Run:

```
switch (config) # ssh server security strict
```

Step 6. Disable HTTP access. Run:

```
switch (config) # no web http enable
```

Step 7. Enable HTTPS strict cyphers. Run:

```
switch (config) # web https ssl ciphers TLS1.2
```

Step 8. Disable router BGP neighbor password configuration. Run:

```
switch (config) # no router bgp <as-number> neighbor <ip-address> password
```

Step 9. Disable router BGP peer group password configuration. Run:

```
switch (config) # no router bgp <as-number> peer-group <peer-group-name> password
```

Step 10. Disable BGP password configuration. Run:

```
switch (config) # no neighbor <ip-address> password
```

Step 11. Disable MD5 password hashing on for users. Run:

```
switch (config) # username <username> password <password>
```

If a necessary prerequisite is not fulfilled the system does not activate secure mode and issues an advisory message accordingly.

Secure mode is not supported on director switch systems.

- ***To activate secure mode:***

```
switch (config) # system secure-mode enable

Warning! Configuration is about to be saved and the system will be reloaded.
  Type 'YES' to confirm the change in secure mode: YES
```

- ***To deactivate secure mode:***

```
switch (config) # no system secure-mode enable

Warning! Configuration is about to be saved and the system will be reloaded.
  Type 'YES' to confirm the change in secure mode: YES
```

- ***To verify secure mode configuration and state:***

```
switch (config)# show system secure-mode

Secure mode configured: yes
Secure mode enabled: yes
  switch (config) #
```

4.13.4 Commands

4.13.4.1 User Accounts

username

username <username> [capability <cap> | disable [login | password] | disconnect | full-name <name> | nopassword | password [0 | 7] <password>]
no username <username> [capability | disable [login | password] | full-name]

Creates a user and sets its capabilities, password and name.
The no form of the command deletes the user configuration.

Syntax Description	username	Specifies a username and creates a user account. New users are created initially with admin privileges but is disabled.
	capability <cap>	Defines user capabilities. • admin - full administrative capabilities • monitor - read only capabilities, can not change the running configuration • unpriv – can only query the most basic information, and cannot take any actions or change any configuration • v_admin – basic administrator capabilities
	disable [login | password]	• Disable - disable this account • Disable login - disable all logins to this account • Disable password - disable login to this account using a local password
	disconnect	Logs out the specified user from the system
	name	Full name of the user
	nopassword	The next login of the user will not require password.
	0 | 7	• 0: specifies a login password in cleartext • 7: specifies a login password in encrypted text
	password	Specifies a password for the user in string form. If [0 | 7] was not specified then the password is in cleartext.
Default	The following usernames are available by default: • admin • monitor • xmladmin • xmluser	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example

	3.4.1100	Updated Example
	3.6.2002	Added “disconnect” parameter

Role

admin

Example

```
switch (config) # username monitor full-name smith
switch (config) # show usernames
USERNAME    FULL NAME               CAPABILITY  ACCOUNT STATUS
USERID      System Administrator    admin       Password set
admin       System Administrator    admin       Password set
monitor     smith                   monitor     Password set (SHA512)
xmladmin    XML Admin User          admin       Password set (SHA512)
xmluser     XML Monitor User        monitor     Password set (SHA512)
switch (config) #
```

Related Commands

show usernames
show users

Notes

- To enable a user account, just set a password on it (or use the command `username <user> nopassword` to enable it with no password required for login)
- Removing a user account does not terminate any current sessions that user has open; it just prevents new sessions from being established
- Encrypted password is useful for the command `show configuration`, since the cleartext password cannot be recovered after it is set

show usernames

	show usernames Displays list of users and their capabilities.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show usernames`
	`USERNAME  FULL NAME  CAPABILITY  ACCOUNT STATUS`
	`USERID  System Administrator  admin  Password set`
	`admin  System Administrator  admin  Password set`
	`monitor  smith  monitor  Password set (SHA512)`
	`xmladmin  XML Admin User  admin  No password required`
	`xmluser  XML Monitor User  monitor  No password required`
	`switch (config) #`

Related Commands	username show users
Notes	

show users

show users [history]

Displays logged in users and related information such as idle time and what host they have connected from.

Syntax Description	history	Displays current and historical sessions.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show users`
	`USERNAME  FULL NAME  LINE  HOST  IDLE`
	`admin  System Administrator  pts/0  172.22.237.174  0d0h34m4s`
	`admin  System Administrator  pts/1  172.30.0.127  1d3h30m49s`
	`admin  System Administrator  pts/3  172.22.237.34  0d0h0m0s`
	`switch (config) #show users history`
	`admin  pts/3 172.22.237.34  Wed Feb  1 11:56  still logged in`
	`admin  pts/3 172.22.237.34  Wed Feb  1 11:42 - 11:46  (00:04)`
	
	`wtmp begins Wed Feb  1 11:38:10 2012`
	`switch (config) #`	

Related Commands	username show usernames	
Notes		

show whoami

	show whoami Displays username and capabilities of user currently logged in.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show whoami`
	`Current user: admin`
	`Capabilities: admin`
	`switch (config) #`

Related Commands	username show usernames show users
Notes	

4.13.4.2 AAA Methods

aaa accounting

aaa accounting changes default stop-only tacacs+
no aaa accounting changes default stop-only tacacs+

Enables logging of system changes to an AAA accounting server.
The no form of the command disables the accounting.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.2.3000 Removed ‘time’ parameter from the command.
Role	admin

Example	`switch (config) # aaa accounting changes default stop-only tacacs+`
	`switch (config) # show aaa`
	`AAA authorization:`
	`  Default User: admin`
	`  Map Order: local-only`
	`Authentication method(s):`
	`  local`
	`  radius`
	`  tacacs+`
	`  ldap`
	`Accounting method(s):`
	`  tacacs+`
	`switch (config) #`

Related Commands	show aaa
Notes	• TACACS+ is presently the only accounting service method supported • Change accounting covers both configuration changes and system actions that are visible under audit logging, however this feature operates independently of audit logging, so it is unaffected by the “logging level audit mgmt” or “configuration audit” commands • Configured TACACS+ servers are contacted in the order in which they appear in the configuration until one accepts the accounting data, or the server list is exhausted • Despite the name of the “stop-only” keyword, which indicates that this feature logs a TACACS+ accounting “stop” message, and in contrast to configuration change accounting, which happens after configuration database changes, system actions are logged when the action is started, not when the action has completed

aaa authentication login

	aaa authentication login default <auth method> [<auth method> [<auth method> [<auth method> [<auth method>]]]] **no aaa authentication login** Sets a sequence of authentication methods. Up to four methods can be configured. The no form of the command resets the configuration to its default.	
Syntax Description	auth-method	• local • radius • tacacs+ • ldap
Default	local	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.7.11xx	Updated notes
Role	admin	

Example	`switch (config) # aaa authentication login default local radius tacacs+ ldap`
	`switch (config) # show aaa`
	`AAA authorization:`
	`  Default User: admin`
	`  Map Order: local-only`
	`Authentication method(s):`
	`  local`
	`  radius`
	`  tacacs+`
	`  ldap`
	`Accounting method(s):`
	`  tacacs+`
	`switch (config) #`	

Related Commands	show aaa	
Notes	The order in which the methods are specified is the order in which the authentication is attempted. It is recommended that “local” is one of the methods selected.	

aaa authentication attempts fail-delay

aaa authentication attempts fail-delay <time>
no aaa authentication attempts fail-delay

Configures delay for a specific period of time after every authentication failure. The no form of the command resets the fail-delay to its default value.

Syntax Description	time	Range: 0-60 seconds
Default	0	
Configuration Mode	config	
History	3.5.0200	
Role	admin	
Example	`switch (config) # aaa authentication attempts fail-delay 1`	
Related Commands	N/A	
Notes		

aaa authentication attempts track

aaa authentication attempts track {downcase | enable}
no aaa authentication attempts track {downcase | enable}

Configure tracking for failed authentication attempts.
The no form of the command clears configuration for tracking authentication failures.

Syntax Description	downcase	Does not convert all usernames to lowercase (for authentication failure tracking purposes only).
	enable	Disables tracking of failed authentication attempts
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # aaa authentication attempts track enable`	
Related Commands	N/A	
Notes	• This is required for the lockout functionality described below, but can also be used on its own for informational purposes. • Disabling tracking does not clear any records of past authentication failures, or the locks in the database. However, it does prevent any updates to this database from being made: no new failures are recorded. It also disables lockout, preventing new lockouts from being recorded and existing lockouts from being enforced.	

aaa authentication attempts lockout

aaa authentication attempts lockout {enable | lock-time | max-fail | unlock-time}
no aaa authentication attempts lockout {enable | lock-time | max-fail | unlock-time}

Configures lockout of accounts based on failed authentication attempts.
The no form of the command clears configuration for lockout of accounts based on failed authentication attempts.

Syntax Description	enable	Enables locking out of user accounts based on authentication failures. This both suspends enforcement of any existing lockouts, and prevents any new lockouts from being recorded. If lockouts are later re-enabled, any lockouts that had been recorded previously resume being enforced; but accounts which have passed the max-fail limit in the meantime are NOT automatically locked at this time. They would be permitted one more attempt, and then locked, because of how the locking is done: lockouts are applied after an authentication failure, if the user has surpassed the threshold at that time. Lockouts only work if tracking is enabled. Enabling lockouts automatically enables tracking. Disabling tracking automatically disables lockouts.
	lock-time	Sets maximum permitted consecutive authentication failures before locking out users. Unlike the “max-fail” setting, this does take effect immediately for all accounts If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time This is not based on the number of consecutive failures, and is therefore divorced from most of the rest of the tally feature, except for the tracking of the last login failure
	max-fail	Sets maximum permitted consecutive authentication failures before locking out users. This setting only impacts what lockouts are imposed while the setting is active; it is not retroactive to previous logins. So if max-fail is disabled or changed, this does not immediately cause any users to be changed from locked to unlocked or vice-versa.
	unlock-time	Enables the auto-unlock of an account after a specified number of seconds if a user account is locked due to authentication failures, counting from the last valid login attempt. Unlike the “max-fail” setting, this does take effect immediately for all accounts. If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time. Careful with disabling the unlock-time, particularly if you have max-fail set to something, and have not overridden the behavior for the admin (i.e. they are subject to lockouts also). If the admin account gets locked out, and there are no other administrators who can aid, the user may be forced to boot single-user and use the pam_tallybyname command-line utility to unlock your account manually. Even if one is careful not to incur this many authentication failures, it makes the system more subject to DOS attacks.

Default	N/A
Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config) # aaa authentication attempts lockout enable`
Related Commands	N/A
Notes	

aaa authentication attempts class-override

aaa authentication attempts class-override {admin [no-lockout] | unknown {no-track | hash-username}}
no aaa authentication attempts class-override {admin | unknown {no-track | hash-username}}

Overrides the global settings for tracking and lockouts for a type of account.
The no form of the command removes this override and lets the admin be handled according to the global settings.

Syntax Description	admin	Overrides the global settings for tracking and lockouts for the admin account. This applies only to the single account with the username “admin”. It does not apply to any other users with administrative privileges.
	no-lockout	Prevents the admin user from being locked out, though the authentication failure history is still tracked (if tracking is enabled overall).
	unknown	Overrides the global settings for tracking and lockouts for unknown accounts. The “unknown” class here contains the following categories: • Real remote usernames which simply failed authentication • Mis-typed remote usernames • Passwords accidentally entered as usernames • Bogus usernames made up as part of an attack on the system
	hash-username	Applies a hash function to the username, and stores the hashed result in lieu of the original.
	no-track	Does not track authentication for such users (which of course also implies no-lockout).
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # aaa authentication attempts class-override admin no-lockout`	
Related Commands	N/A	
Notes		

aaa authentication attempts reset

aaa authentication attempts reset {all | user <username>} [{no-clear-history | no-unlock}]

Clears the authentication history for and/or unlocks specified users.

Syntax Description	all	Applies function to all users.
	user	Applies function to specified user.
	no-clear-history	Leaves the history of login failures but unlocks the account.
	no-unlock	Leaves the account locked but clears the history of login failures.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # aaa authentication attempts reset user admin all`	
Related Commands	N/A	
Notes		

clear aaa authentication attempts

clear aaa authentication attempts {all | user <username>} [no-clear-history | no-unlock]

Clears the authentication history for and/or unlocks specified users

Syntax Description	all	Applies function to all users.
	user	Applies function to specified user.
	no-clear-history	Clears the history of login failures.
	no-unlock	Unlocks the account.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # aaa authentication attempts reset user admin no-clear-history`	
Related Commands	N/A	
Notes		

aaa authorization

aaa authorization map [default-user <username> | order <policy> | fallback]
no aaa authorization map [default-user | order | fallback]

Sets the mapping permissions of a user in case a remote authentication is done. The no form of the command resets the attributes to default.

Syntax Description	username	Specifies what local account the authenticated user will be logged on as when a user is authenticated (via RADIUS or TACACS+ or LDAP) and does not have a local account. If the username is local, this mapping is ignored.
	order <policy>	Sets the user mapping behavior when authenticating users via RADIUS or TACACS+ or LDAP to one of three choices. The order determines how the remote user mapping behaves. If the authenticated username is valid locally, no mapping is performed. The setting has the following three possible behaviors: • local-only – maps all remote users to the user specified by the "aaa authorization map default-user <user name>" command. Any vendor attributes received by an authentication server are ignored. • remote-first – if a local-user mapping attribute is returned and it is a valid local username, it maps the authenticated user to the local user specified in the attribute. Otherwise, it uses the user specified by the default-user command. • remote-only – maps a remote authenticated user if the authentication server sends a local-user mapping attribute. If the attribute does not specify a valid local user, no further mapping is tried.
	fallback	Sets the authenticating fallback behavior via RADIUS or TACACS+ or LDAP. This option attempts to authenticate username through the next authentication method listed in case of an error. • server-err – performs fallback if an error occurs while connecting to remote AAA server (e.g. server is down, not responding, etc)
Default	Default user – admin Map order – remote-first Order fallback – server-err	
Configuration Mode	config	
History	3.1.0000	
	3.7.1000	Added "fallback" parameter

	3.7.1000	Updated syntax
Role	admin	
Example	`switch (config) # aaa authorization map default-user admin` `switch (config) #`	
Related Commands	show aaa username	
Notes	• If, for example, the user is locally defined to have admin permission, but in a remote server such as RADIUS the user is authenticated as monitor and the order is remote-first, then the user is given monitor permissions. • If AAA authorization order policy is configured to remote-only, then when upgrading to 3.4.3000 or later from an older Onyx version, this policy is changed to remote-first. • The user must be careful when disabling AAA authorization map fallback server-err, because if the remote server stops working then the user may lock themselves out.	

show aaa

	show aaa Displays the AAA configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.0020 Example updated
Role	admin

Example	`switch (config) # show aaa`
	`AAA authorization:`
	`  Default User: admin`
	`  Map Order: remote-first`
	`  Fallback on server-err: yes`
	`Authentication method(s):`
	`  local`
	`Accounting method(s):`
	`  tacacs+`
	`switch (config) #`

Related Commands	aaa accounting aaa authentication aaa authorization show aaa show usernames username
Notes	

show aaa authentication attempts

show aaa authentication attempts [configured | status user <username>]]

Shows the current authentication, authorization and accounting settings.

Syntax Description	authentication attempts	Displays configuration and history of authentication failures.
	configured	Displays configuration of authentication failure tracking.
	status user	Displays status of authentication failure tracking and lockouts for specific user.
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000	
	3.5.0200	Updated Example
Role	admin	

Example

```
switch (config) # show aaa authentication attempts
Configuration for authentication failure tracking and locking:
   Track authentication failures:                    yes
   Lock accounts based on authentication failures:   yes
   Override treatment of 'admin' user:               (none)
   Override treatment of unknown usernames:          hash-usernames
   Convert usernames to lowercase for tracking:      no
   Delay after each auth failure (fail delay):       none


Configuration for lockouts based on authentication failures:
   Lock account after consecutive auth failures:    5
   Allow retry on locked accounts (unlock time):    after 15 second(s)
   Temp lock after each auth failure (lock time):   none

Username                              Known  Locked  Failures  Last fail time        Last fail from
--------                              -----  ------  --------  --------------        --------------
0Q72B43EHBKT8CB5AF5PGRX3U3B3TUL4CYJP93N(*) no     no         1     2012/08/20 14:29:19   ttyS0

(*) Hashed for security reasons
switch-627d3c [standalone: master] (config) #
switch (config) #
```

Related Commands	N/A
Notes	

4.13.4.3 RADIUS

radius-server

radius-server {key <secret>| retransmit <retries> | timeout <seconds>}
no radius-server {key | retransmit | timeout}

Sets global RADIUS server attributes.
The no form of the command resets the attributes to their default values.

Syntax Description	secret	Sets a secret key (shared hidden text string), known to the system and to the RADIUS server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry (1-60).
Default	3 seconds, 1 retry	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # radius-server retransmit 3`	
Related Commands	aaa authorization radius-server host show radius	
Notes	Each RADIUS server can override those global parameters using the command “radius-server host”.	

radius-server host

radius-server host <IP address> [enable | auth-port <port> | key <secret> | prompt-key | retransmit <retries> | timeout <seconds>]
no radius-server host <IP address> [auth-port | enable]

Configures RADIUS server attributes.
The no form of the command resets the attributes to their default values and deletes the RADIUS server.

Syntax Description	IP address	RADIUS server IP address
	enable	Administrative enable of the RADIUS server
	auth-port	Configures authentication port to use with this RADIUS server
	port	RADIUS server UDP port number
	key	Configures shared secret to use with this RADIUS server
	prompt-key	Prompt for key, rather than entering on command line
	retransmit	Configures retransmit count to use with this RADIUS server
	retries	Number of retries (0-5) before exhausting from the authentication
	timeout	Configures timeout between each try
	seconds	Timeout in seconds between each retry (1-60)
Default	3 seconds, 1 retry Default UDP port is 1812	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # radius-server host 40.40.40.40`	
Related Commands	aaa authorization radius-server show radius	
Notes	• RADIUS servers are tried in the order they are configured • If you do not specify a parameter for this configured RADIUS server, the configuration will be taken from the global RADIUS server configuration. Refer to “radius-server” command.	

show radius

	show radius Displays RADIUS configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show radius`
	`RADIUS defaults:`
	`  Key       : ********`
	`  Timeout   : 3`
	`  Retransmit: 1`
	
	`RADIUS servers:`
	`  1.1.1.1:1812:`
	`    Enabled   : yes`
	`    Key       : ********`
	`    Timeout   : 3 (default)`
	`    Retransmit: 1 (default)`

Related Commands	aaa authorization radius-server radius-server host
Notes	

4.13.4.4 TACACS+

tacacs-server

tacacs-server {key <secret>| retransmit <retries> | timeout <seconds>}
no tacacs-server {key | retransmit | timeout}

Sets global TACACS+ server attributes.
The no form of the command resets the attributes to default values.

Syntax Description	secret	Set a secret key (shared hidden text string), known to the system and to the TACACS+ server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry (1-60).
Default	3 seconds, 1 retry	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # tacacs-server retransmit 3`	
Related Commands	aaa authorization show radius show tacacs tacacs-server host	
Notes	Each TACACS+ server can override those global parameters using the command “tacacs-server host”.	

tacacs-server host

tacacs-server host <IP address> {enable | auth-port <port> | auth-type <type> | key <secret> | prompt-key | retransmit <retries> | timeout <seconds>}
no tacacs-server host <IP address> {enable | auth-port}

Configures TACACS+ server attributes.
The no form of the command resets the attributes to their default values and deletes the TACACS+ server.

Syntax Description	IP address	TACACS+ server IP address
	enable	Administrative enable for the TACACS+ server
	auth-port	Configures authentication port to use with this TACACS+ server
	port	TACACS+ server UDP port number
	auth-type	Configures authentication type to use with this TACACS+ server
	type	Authentication type. Possible values are: • ASCII • PAP (Password Authentication Protocol)
	key	Configures shared secret to use with this TACACS+ server
	secret	Sets a secret key (shared hidden text string), known to the system and to the TACACS+ server
	prompt-key	Prompts for key, rather than entering key on command line
	retransmit	Configures retransmit count to use with this TACACS+ server
	retries	Number of retries (0-5) before exhausting from the authentication
	timeout	Configures timeout to use with this TACACS+ server
	seconds	Timeout in seconds between each retry (1-60)
Default	3 seconds, 1 retry Default TCP port is 49 Default auth-type is PAP	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # tacacs-server host 40.40.40.40`
Related Commands	aaa authorization show tacacs tacacs-server
Notes	• TACACS+ servers are tried in the order they are configured • A PAP auth-type similar to an ASCII login, except that the username and password arrive at the network access server in a PAP protocol packet instead of being typed in by the user, so the user is not prompted • If the user does not specify a parameter for this configured TACACS+ server, the configuration will be taken from the global TACACS+ server configuration. Refer to “tacacs-server” command.

show tacacs

	show tacacs Displays TACACS+ configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show tacacs`
	`TACACS+ defaults:`
	`  Key       :  ********`
	`  Timeout   :  3`
	`  Retransmit:  1`
	
	`TACACS+ servers:`
	`  1.1.1.1:49:`
	`    Enabled   : yes`
	`    Auth Type : pap`
	`    Key       : ********`
	`    Timeout   : 3 (default)`
	`    Retransmit: 1 (default)`

Related Commands	aaa authorization tacacs-server tacacs-server host
Notes	

4.13.4.5 LDAP

ldap base-dn

ldap base-dn <string>
no ldap base-dn

Sets the base distinguished name (location) of the user information in the schema of the LDAP server.
The no form of the command resets the attribute to its default values.

Syntax Description	string	A case-sensitive string that specifies the location in the LDAP hierarchy where the server should begin searching when it receives an authorization request. For example: “ou=users,dc=example,dc=com”, with no spaces. when: ou - Organizational unit dc - Domain component cn - Common name sn - Surname
Default	ou=users,dc=example,dc=com	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap base-dn ou=department,dc=example,dc=com`	
Related Commands	show ldap	
Notes		

ldap bind-dn/bind-password

ldap {bind-dn | bind-password} <string>
no ldap {bind-dn | bind-password}

Gives the distinguished name or password to bind to on the LDAP server. This can be left empty for anonymous login (the default).
The no form of the command resets the attribute to its default values.

Syntax Description	string	A case-sensitive string that specifies distinguished name or password to bind to on the LDAP server.
Default	“”	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap bind-dn my-dn` `switch (config) # ldap bind-password my-password`	
Related Commands	show ldap	
Notes	For anonymous login, bind-dn and bind-password should be empty strings “”.	

ldap group-attribute/group-dn

ldap {group-attribute {<group-att> |member | uniqueMember} | group-dn <group-dn>}
no ldap {group-attribute | group-dn}

Sets the distinguished name or attribute name of a group on the LDAP server. The no form of the command resets the attribute to its default values.

Syntax Description	group-att	Specifies a custom attribute name.
	member	groupOfNames or group membership attribute.
	uniqueMember	groupOfUniqueNames membership attribute.
	group-dn	DN of group required for authorization.
Default	group-att: member group-dn: “”	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap group-attribute member` `switch (config) # ldap group-dn my-group-dn`	
Related Commands	show ldap	
Notes	• The user’s distinguished name must be listed as one of the values of this attribute, or the user will not be authorized to log in • After login authentication, if the group-dn is set, a user must be a member of this group or the user will not be authorized to log in. If the group is not set (“” - the default) no authorization checks are done.	

ldap host

ldap host <IP Address> [order <number> last]
no ldap host <IP Address>

Adds an LDAP server to the set of servers used for authentication.
The no form of the command deletes the LDAP host.

Syntax Description	IP Address	IPv4 or IPv6 address.
	number	The order of the LDAP server.
	last	The LDAP server will be added in the last location.
Default	No hosts configured	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap host 10.10.10.10`	
Related Commands	show aaa show ldap	
Notes	• The system will select the LDAP host to try according to its order • New servers are by default added at the end of the list of servers	

ldap hostname-check enable

	ldap hostname-check enable **no ldap hostname-check enable** Enables LDAP hostname check. The no form of the command disables LDAP hostname check.
Syntax Description	N/A
Default	No hosts configured
Configuration Mode	config
History	3.6.8008
Role	admin
Example	`switch (config) # ldap hostname-check enable`
Related Commands	show aaa show ldap
Notes	

ldap login-attribute

ldap login-attribute {<string> | uid | sAMAccountName}
no ldap login-attribute

Sets the attribute name which contains the login name of the user.
The no form of the command resets this attribute to its default.

Syntax Description	string	Custom attribute name.
	uid	LDAP login name is taken from the user login user-name.
	sAMAccountName	SAM Account name, active directory login name.
Default	sAMAccountName	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap login-attribute uid`	
Related Commands	show aaa show ldap	
Notes		

ldap port

ldap port <port>
no ldap port

Sets the TCP port on the LDAP server to connect to for authentication.
The no form of the command resets this attribute to its default value.

Syntax Description	port	TCP port number.
Default	389	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap port 1111`	
Related Commands	show aaa show ldap	
Notes		

ldap referrals

	ldap referrals **no ldap referrals** Enables LDAP referrals. The no form of the command disables LDAP referrals.
Syntax Description	N/A
Default	LDAP referrals are enabled
Configuration Mode	config
History	3.1.0000 3.4.0000 Updated Example
Role	admin
Example	`switch (config) # no ldap referrals`
Related Commands	show aaa show ldap
Notes	Referral is the process by which an LDAP server, instead of returning a result, will return a referral (a reference) to another LDAP server which may contain further information.

ldap scope

	ldap scope <scope> **no ldap scope** Specifies the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request. The no form of the command resets the attribute to its default value.	
Syntax Description	scope	• one-level - searches the immediate children of the base dn • subtree - searches at the base DN and all its children
Default	subtree	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap scope subtree`	
Related Commands	show aaa show ldap	
Notes		

ldap ssl

ldap ssl {ca-list <options> | cert-verify | ciphers {all | TLS1.2} | crl-check {enable | file fetch <path>} | mode <mode> | port <port-number>}
no ldap ssl {cert-verify | ciphers | crl-check enable | mode | port}

Sets SSL parameter for LDAP.
The no form of the command resets the attribute to its default value.

Syntax Description	options	This command specifies the list of supplemental certificates of authority (CAs) from the certificate configuration database that is to be used by LDAP for authentication of servers when in TLS or SSL mode. The options are: • default-ca-list - uses default supplemental CA certificate list • none - no supplemental list, uses the built-in one only CA certificates are ignored if "ldap ssl mode" is not configured as either "tls" or "ssl", or if "no ldap ssl cert-verify" is configured. The default-ca-list is empty in the factory default configuration. Use the command: "crypto certificate ca-list default-ca-list name" to add trusted certificates to that list. The "default-ca-list" option requires LDAP to consult the system's configured global default CA-list for supplemental certificates.
	cert-verify	Enables verification of SSL/TLS server certificates. This may be required if the server's certificate is self-signed, or does not match the name of the server.
	ciphers {all | TLS1.2}	Sets SSL mode to be used.
	crl-check enable	Enables LDAP CRL check
	crl-check file fetch	Fetches CRL from remote server. CRL must be a valid PEM file unless a proper message shown. Supported formats: SCP, HTTP, HTTPS, FTP, and FTPS.
	mode	Sets the security mode for connections to the LDAP server. • none – requests no encryption for the LDAP connection • ssl – the SSL-port configuration is used, an SSL connection is made before LDAP requests are sent (LDAP over SSL) • start-tls – the normal LDAP port is used, an LDAP connection is initiated, and then TLS is started on this existing connection
	port-number	Sets the port on the LDAP server to connect to for authentication when the SSL security mode is enabled (LDAP over SSL).
Default	cert-verify: enabled mode: none (LDAP SSL is not activated) port-number: 636 ciphers: all	
Configuration Mode	config	

History	3.1.0000	
	3.2.3000	Added ca-list argument.
	3.4.0000	Added “ssl ciphers” parameter and updated Example
	3.6.8008	Added the parameter “crl-check”
Role	admin	

Example	`switch (config) # ldap ssl crl-check file fetch scp://root:pass@1.1.1.1/etc/pki/crl.pem`
	
	`100.0%`
	`[##########################################################################]`	

Related Commands	show aaa show ldap	
Notes	• If available, the TLS mode is recommended, as it is standardized, and may also be of higher security • The port number is used only for SSL mode. In case the mode is TLS, the LDAP port number will be used.	

ldap timeout

ldap {timeout-bind | timeout-search} <seconds>
no ldap {timeout-bind | timeout-search}

Sets a global communication timeout in seconds for all LDAP servers to specify the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request.
The no form of the command resets the attribute to its default value.

Syntax Description	timeout-bind	Sets the global LDAP bind timeout for all LDAP servers.
	timeout-search	Sets the global LDAP search timeout for all LDAP servers.
	seconds	Range: 1-60 seconds.
Default	5 seconds	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap timeout-bind 10`	
Related Commands	show aaa show ldap	
Notes		

ldap version

ldap version <version>
no ldap version

Sets the LDAP version.
The no form of the command resets the attribute to its default value.

Syntax Description	version	Sets the LDAP version. Values: 2 and 3.
Default	3	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	`switch (config) # ldap version 3`	
Related Commands	show aaa show ldap	
Notes		

show ldap

	show ldap Displays LDAP configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.0000 Updated Example 3.6.8008 Updated Example
Role	admin

Example	`switch (config) # show ldap`
	
	`User base DN : ou=users,dc=example,dc=com`
	`User search scope : subtree`
	`Login attribute : sAMAccountName`
	`Bind DN :`
	`Bind password : ********`
	`Group base DN :`
	`Group attribute : member`
	`LDAP version : 3`
	`Referrals : yes`
	`Server port : 389`
	`Search Timeout : 5`
	`Bind Timeout : 5`
	`Server Hostname check: no`
	`SSL mode : none`
	`Server SSL port : 636 (not active)`
	`SSL ciphers : all (not active)`
	`SSL cert verify : yes`
	`SSL ca-list : default-ca-list`
	`SSL CRL check : no`

Related Commands	show aaa show ldap
Notes	

show ldap crl

	show ldap crl Displays current CRL configured by the user.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Role	admin

Example	`switch (config) # show ldap crl`
	`-----BEGIN CERTIFICATE-----`
	`MIIDVzCSd……`
	`-----END CERTIFICATE-----`

Related Commands	show aaa show ldap
Notes	

4.13.4.6 System Secure Mode

system secure-mode enable

system secure-mode enable
no system secure-mode enable

Enables secure mode on the switch.
The no form of the command disables secure mode.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.0200
Role	admin

Example	`switch (config) # system secure-mode enable`
	
	`Warning! Configuration is about to be saved and the system will be reloaded.`
	`Type 'YES' to confirm the change in secure mode: YES`

Related Commands	user <username> password <password> ssh server min-version ssh server security strict snmp-server user no neighbor <ip-address> password ntp server disable ntp server keyID router bgp neighbor password router bgp peer-group password
Notes	Before enabling secure mode, the command performs the following configuration checks: • NTP Key ID cannot be MD5 when secure mode is enabled • SSH min-version cannot be 1 when enabling secure mode • SSH security must be set to strict security • SNMPv3 user auth cannot be md5 when enabling secure mode • SNMPv3 user priv cannot be des when enabling secure mode • SNMPv3 trap auth cannot be md5 when enabling secure mode • SNMPv3 trap priv cannot be des when enabling secure mode • Router BGP neighbor password cannot be set when enabling secure mode • Router BGP peer-group password cannot be set when enabling with secure mode • User password hash cannot be MD5 when secure mode is enabled Only if the check passes, secure mode is enabled on the switch system.

show system secure-mode

	show system secure-mode Displays the security mode of the switch system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2300
Role	admin

Example	`switch (config) # show system secure-mode`
	
	`Secure mode configured: yes`
	`Secure mode enabled : yes`
	`switch (config) #`

Related Commands	system secure-mode enable
Notes	“Secure mode configuration” describes the user configuration “Secure mode enabled” describes the system state

4.14 Cryptographic (X.509, IPSec) and Encryption

This chapter contains commands for configuring, generating and modifying x.509 certificates used in the system. Certificates are used for creating a trusted SSL connection to the system.

Crypto commands also cover IPSec configuration commands used for establishing a secure connection between hosts over IP layer which is useful for transferring sensitive information.

4.14.1 System File Encryption

This feature encrypts all sensitive data on Mellanox systems including logs certificates, keys, etc.

➢ ***To activate encryption on the switch:***

Step 1. Enable encryption and configure key location as USB (if you are using a USB device). Run:

```
switch (config)# crypto encrypt-data key-location usb key mypassword

Warning! All sensitive files are about to be encrypted
- System will perform reset factory, configuration files will be preserved
- System will be rebooted
- Active configuration will be preserved
- Do not power-off, wait for the system to boot

Type 'YES' to confirm this action: YES
```

IMPORTANT NOTE
Encryption and decryption perform “reset factory keep-config” on the switch system once configured. This means that sysdumps, logs, and images are deleted.

The key may be saved locally as well by using the parameter “local” instead of “usb” but that configuration is less secure.

Step 2. After the system reboots, verify configuration. Run:

```
switch (config)# show crypto encrypt-data
Sensitive files encryption:
   Status:          enabled
   Key location:    usb
   Cipher:          aes256
```

Once encryption is enabled, reverting back to an older version while encrypted is not possible. The command “no crypto encrypt-data” must be run before attempting to downgrade to an older Onyx version.

If encryption is enabled, upgrading to a new Onyx version maintains the encryption configuration.

4.14.1.1 Commands

crypto encrypt-data

crypto encrypt-data key-location <local | usb> key <password>
no crypto encrypt-data

Enables and configures system file encryption.
The no form of the command decrypts sensitive information on the system.

Syntax Description	key-location	Configures where to store the encryption key: • local – Stores the key locally • usb – Stores the key on a USB device
	key	Configures a key
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
Role	admin	

Example	`switch (config)# crypto encrypt-data key-location usb key mypassword`
	
	`Warning! All sensitive files are about to be encrypted`
	`- System will perform reset factory, configuration files will be preserved`
	`- System will be rebooted`
	`- Do not power-off, wait for the system to boot`
	
	`Type 'YES' to confirm this action: YES`	

Related Commands		
Notes	• It is recommended to store the encryption password on a USB device rather than locally • Enabling encryption may slightly slow system performance • If the key is stored on the USB, it must be plugged into the switch in order for the switch to boot. After the switch has booted, the USB key is no longer required and, for security purposes, it is recommended to remove it after running “usb eject”. The USB key may be needed again if the switch is rebooted or if the switch needs to be decrypted.	

crypto ipsec ike

crypto ipsec ike {clear sa [peer {any | <IPv4 or IPv6 address>} local <IPv4 or IPv6 address>] | restart}

Manage the IKE (ISAKMP) process or database state

Syntax Description	clear	Clears IKE (ISAKMP) peering state
	sa	Clears IKE generated ISAKMP and IPSec security associations (remote peers are affected)
	peer	Clears security associations for the specified IKE peer (remote peers are affected) all – clears security associations for all IKE peerings with a specific local address (remote peers are affected) IPv4 or IPv6 address – clears security associations for specific IKE peering with a specific local address (remote peers are affected)
	IPv4 or IPv6 address	Clears security associations for the specified IKE peering (remote peer is affected)
	local	Clear security associations for the specified/all IKE peering (remote peer is affected)
	restart	Restarts the IKE (ISAKMP) daemon (clears all IKE state, peers may be affected)
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config)# crypto ipsec ike restart` `switch (config)#`	
Related Commands	N/A	
Notes		

crypto ipsec peer local

crypto ipsec peer <IPv4 or IPv6 address> local <IPv4 or IPv6 address> {enable | keying {ike [auth {hmac-md5 | hmac-sha1 | hmac-sha256 | null} | dh-group | disable | encrypt | exchange-mode | lifetime | local | mode | peer-identity | pfs-group | preshared-key | prompt-preshared-key | transform-set] | manual [auth | disable | encrypt | local-spi | mode | remote-spi]}}

Configures ipsec in the system.

Syntax Description	enable	Enables IPSec peering.
	ike	Configures IPSec peering using IKE ISAKMP to manage SA keys. It has the following optional parameters: • auth: Configures the authentication algorithm for IPSec peering • dh-group: Configures the phase1 Diffie-Hellman group proposed for secure IKE key exchange • disable: Configures this IPSec peering administratively disabled • encrypt: Configures the encryption algorithm for IPSec peering • exchange-mode: Configures the IKE key exchange mode to propose for peering • lifetime: Configures the SA lifetime to propose for this IPSec peering • local-identity: Configures the ISAKMP payload identification value to send as local endpoint's identity • mode: Configures the peering mode for this IPSec peering • peer-identity: Configures the identification value to match against the peer's ISAKMP payload identification • pfs-group: Configures the phase2 PFS (Perfect Forwarding Secrecy) group to propose for Diffie-Hellman exchange for this IPSec peering • preshared-key: Configures the IKE pre-shared key for the IPSec peering • prompt-preshared-key: Prompts for the pre-shared key, rather than entering it on the command line • transform-set: Configures transform proposal parameters
	keying	Configures key management for this IPSec peering: • auth: Configures the authentication algorithm for this IPSec peering • disable: Configures this IPSec peering administratively disabled • encrypt: Configures the encryption algorithm for this IPSec peering • local-spi: Configures the local SPI for this manual IPSec peering • mode: Configures the peering mode for this IPSec peering • remote-spi: Configures the remote SPI for this manual IPSec peering
	manual	Configures IPSec peering using manual keys.
Default	N/A	

Configuration Mode	config
History	3.2.3000
Role	admin
Example	`switch (config)# crypto ipsec peer 10.10.10.10 local 10.7.34.139 enable` `switch (config)#`
Related Commands	N/A
Notes	

crypto certificate ca-list

crypto certificate ca-list [default-ca-list name {<cert-name> | system-self-signed}]
no crypto certificate ca-list [default-ca-list name {<cert-name> | system-self-signed}]

Adds the specified CA certificate to the default CA certificate list.
The no form of the command removes the certificate from the default CA certificate list.

Syntax Description	cert-name	The name of the certificate.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # crypto certificate default-cert name test`	
Related Commands	N/A	
Notes	• Two certificates with the same subject and issuer fields cannot both be placed onto the CA list • The no form of the command does not delete the certificate from the certificate database • Unless specified otherwise, applications that use CA certificates will still consult the well-known certificate bundle before looking at the default-ca-list	

crypto certificate default-cert

crypto certificate default-cert name {<cert-name> | system-self-signed}
no crypto certificate default-cert name {<cert-name> | system-self-signed}

Designates the named certificate as the global default certificate role for authentication of this system to clients.
The no form of the command reverts the default-cert name to “system-self-signed” (the “cert-name” value is optional and ignored).

Syntax Description	cert-name	The name of the certificate.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # crypto certificate default-cert name test`	
Related Commands	N/A	
Notes	• A certificate must already be defined before it can be configured in the default-cert role • If the named default-cert is deleted from the database, the default-cert automatically becomes reconfigured to the factory default, the “system-self-signed” certificate	

crypto certificate generation

crypto certificate generation default {country-code | days-valid | email-addr | hash-algorithm {sha1 | sha256} | key-size-bits | locality | org-unit | organization | state-or-prov}

Configures default values for certificate generation.

Syntax Description	country-code	Configures the default certificate value for country code with a two-alphanumeric-character code or -- for none.
	days-valid	Configures the default certificate valid days. Default value: 365 days.
	email-addr	Configures the default certificate value for email address.
	hash-algorithm {sha1 | sha256}	Configures the default certificate hashing algorithm.
	key-size-bits	Configures the default certificate value for private key size. (Private key length in bits – at least 1024, but 2048 is strongly recommended.)
	locality	Configures the default certificate value for locality.
	org-unit	Configures the default certificate value for organizational unit.
	organization	Configures the default certificate value for the organization name.
	state-or-prov	Configures the default certificate value for state or province.
Default	N/A	
Configuration Mode	config	
History	3.2.1000	
	3.3.4350	Added “hash-algorithm” parameter
	3.6.4000	Added “days-valid” parameter
Role	admin	
Example	`switch (config) # crypto certificate generation default hash-algorithm sha256`	
Related Commands	N/A	
Notes	The default hashing algorithm used is sha1.	

crypto certificate name

crypto certificate name {<cert-name> | system-self-signed} {comment <new comment> | generate self-signed [comment <cert-comment> | common-name <domain> | country-code <code> | days-valid <days> | email-addr <address> | hash-algorithm {sha1 | sha256} | key-size-bits <bits> | locality <name> | org-unit <name> | organization <name> | serial-num <number> | state-or-prov <name>]} | private-key pem <PEM string> | prompt-private-key | public-cert [comment <comment string> | pem <PEM string>] | regenerate days-valid <days> | rename <new name>}
no crypto certificate name <cert-name>

Configures default values for certificate generation.
The no form of the command clears/deletes certain certificate settings.

Syntax Description	cert-name	Unique name by which the certificate is identified.
	comment	Specifies a certificate comment.
	generate self-signed	Generates certificates. This option has the following parameters which may be entered sequentially in any order: • comment: Specifies a certificate comment (free string) • common-name: Specifies the common name of the issuer and subject (e.g. a domain name) • country-code: Specifies the country codwo-alpha-numeric-character country code, or "--" for none) • days-valid: Specifies the number of days the certificate is valid • email-addr: Specifies the email address • hash-algorithm: Specifies the hashing function used for signature algorithm. Default value is SHA256. • key-size-bits: Specifies the size of the private key in bits (private key length in bits - at least 1024 but 2048 is strongly recommended) • locality: Specifies the locality name • org-unit: Specifies the organizational unit name • organization: Specifies the organization name • serial-num: Specifies the serial number for the certificate (a lower-case hexadecimal serial number prefixed with "0x") • state-or-prov: Specifies the state or province name
	private-key pem	Specifies certificate contents in PEM format.
	prompt-private-key	Prompts for certificate private key with secure echo.
	public-cert	Installs a certificate.
	regenerate	Regenerates the named certificate using configured certificate generation default values for the specified validity period
	rename	Renames the certificate.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
	3.3.4402	Added "hash-algorithm" parameter
	3.6.4000	Added "hash-algorithm" parameter
Role	admin	
Example	`switch (config) # crypto certificate name system-self-signed generate self-signed hash-algorithm sha256`	

Related Commands	N/A
Notes	

crypto certificate system-self-signed

crypto certificate system-self-signed regenerate [days-valid <days>]

Configures default values for certificate generation.

Syntax Description	days-valid	Specifies the number of days the certificate is valid
Default	N/A	
Configuration Mode	config	
History	3.2.1000	
Role	admin	
Example	`switch (config) # crypto certificate system-self-signed regenerate days-valid 3`	
Related Commands	N/A	
Notes		

show crypto certificate

show crypto certificate [detail | public-pem | default-cert [detail | public-pem] | [name <cert-name> [detail | public-pem] | ca-list [default-ca-list]]

Displays information about all certificates in the certificate database.

Syntax Description	ca-list	Displays the list of supplemental certificates configured for the global default system CA certificate role.
	default-ca-list	Displays information about the currently configured default certificates of the CA list.
	default-cert	Displays information about the currently configured default certificate.
	detail	Displays all attributes related to the certificate.
	name	Displays information about the certificate specified.
	public-pem	Displays the uninterpreted public certificate as a PEM formatted data string
Default	N/A	
Configuration Mode	config	
History	3.2.1000	
Role	admin	

Example

```
switch (config)# show crypto certificate
Certificate with name 'system-self-signed' (default-cert)
    Comment:                        system-generated self-signed certif-
icate
    Private Key:                    present
    Serial Number:                  0x546c935511bcafc21ac0e8249fbe0844
    SHA-1 Fingerprint:              fe6df38dd26801971cb2d44f62d-
be492b6063c5f

    Validity:
        Starts:                     2012/12/02 13:45:05
        Expires:                    2013/12/02 13:45:05

    Subject:
        Common Name:                IBM-DEV-Bay4
        Country:                    IS
        State or Province:
        Locality:
        Organization:
        Organizational Unit:
        E-mail Address:

    Issuer:
        Common Name:                IBM-DEV-Bay4
        Country:                    IS
        State or Province:
        Locality:
        Organization:
        Organizational Unit:
        E-mail Address:
switch (config)#
```

Related Commands N/A

Notes

show crypto encrypt-data

	show encrypt-data Displays sensitive data encryption information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.1002
Role	admin

Example	`switch (config)# show crypto encrypt-data`
	`Sensitive files encryption:`
	`  Status:  enabled`
	`  Key location:  usb`
	`  Cipher:  aes256`
	`switch (config)#`

Related Commands	N/A
Notes	

show crypto ipsec

show crypto ipsec [brief | configured | ike | policy | sa]

Displays information ipsec configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.1000
Role	admin

Example	`switch (config)# show crypto ipsec`
	`IPSec Summary`
	`-------------`
	`Crypto IKE is using pluto (Openswan) daemon.`
	`Daemon process state is stopped.`
	
	`No IPSec peers configured.`
	
	`IPSec IKE Peering State`
	`----------------------`
	`Crypto IKE is using pluto (Openswan) daemon.`
	`Daemon process state is stopped.`
	
	`No active IPSec IKE peers.`
	
	`IPSec Policy State`
	`------------------`
	`No active IPSec policies.`
	
	`IPSec Security Association State`
	`-------------------------------`
	`No active IPSec security associations.`
	`switch (config)#`

Related Commands	N/A
Notes	

4.15 Scheduled Jobs

Use the commands in this section to manage and schedule the execution of jobs

4.15.1 Commands

job

job <job ID>
no job <job ID>

Creates a job.
The no form of the command deletes the job.

Syntax Description	job ID	An integer.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # job 100` `switch (config job 100) #`	
Related Commands	show jobs	
Notes	Job state is lost on reboot.	

command

command <sequence #> | <command>
no command <sequence #>

Adds a CLI command to the job.
The no form of the command deletes the command from the job.

Syntax Description	sequence #	An integer that controls the order the command is executed relative to other commands in this job. The commands are executed in an ascending order.
	command	A CLI command.
Default	N/A	
Configuration Mode	config job	
History	3.1.0000	
Role	admin	

Example	`switch (config)# job 100`
	`switch (config job 100) # command 10 “show power”`
	`switch (config job 100) #`	

Related Commands	show jobs	
Notes	• The command must be defined with inverted commas (“”) • The command must be added as it was executed from the “config” mode. For example, in order to change the interface description you need to add the command: “interface <type> <number> description my-description”.	

comment

comment <comment>
no comment

Adds a comment to the job.
The no form of the command deletes the comment.

Syntax Description	comment	The comment to be added (string).
Default	“”	
Configuration Mode	config job	
History	3.1.0000	
Role	admin	

Example	`switch (config)# job 100`
	`switch (config job 100) # comment Job_for_example`
	`switch (config job 100) #`	

Related Commands	show jobs	
Notes		

enable

	enable **no enable** Enables the specified job. The no form of the command disables the specified job.
Syntax Description	N/A
Default	N/A
Configuration Mode	config job
History	3.1.0000
Role	admin

Example	`switch (config)# job 100`
	`switch (config job 100) # enable`
	`switch (config job 100) #`

Related Commands	show jobs
Notes	If a job is disabled, it will not be executed automatically according to its schedule; nor can it be executed manually.

execute

	execute Forces an immediate execution of the job.
Syntax Description	N/A
Default	N/A
Configuration Mode	config job
History	3.1.0000
Role	admin

Example	`switch (config)# job 100`
	`switch (config job 100) # execute`
	`switch (config job 100) #`

Related Commands	show jobs
Notes	• The job timer (if set) is not canceled and the job state is not changed: i.e. the time of the next automatic execution is not affected • The job will not be run if not currently enabled

fail-continue

	fail-continue **no fail-continue** Continues the job execution regardless of any job failures. The no form of the command returns fail-continue to its default.
Syntax Description	N/A
Default	A job will halt execution as soon as any of its commands fails
Configuration Mode	config job
History	3.1.0000
Role	admin

Example	`switch (config)# job 100`
	`switch (config job 100) # fail-continue`
	`switch (config job 100) #`

Related Commands	show jobs
Notes	

name

name <job name>
no name

Configures a name for this job.
The no form of the command resets the name to its default.

Syntax Description	name	Specifies a name for the job (string).
Default	“”.	
Configuration Mode	config job	
History	3.1.0000	
Role	admin	

Example	`switch (config)# job 100`
	`switch (config job 100) # name my-job`
	`switch (config job 100) #`	

Related Commands	show jobs	
Notes		

schedule type

schedule type <recurrence type>
no schedule type

Sets the type of schedule the job will automatically execute on.
The no form of the command resets the schedule type to its default.

Syntax Description	recurrence type	The available schedule types are: • daily - the job is executed every day at a specified time • weekly - the job is executed on a weekly basis • monthly - the job is executed every month on a specified day of the month • once - the job is executed once at a single specified date and time • periodic - the job is executed on a specified fixed time interval, starting from a fixed point in time.
Default	once	
Configuration Mode	config job	
History	3.1.0000	
Role	admin	

Example	`switch (config)# job 100`
	`switch (config job 100) # schedule type once`
	`switch (config job 100) #`	

Related Commands	show jobs	
Notes	A schedule type is essentially a structure for specifying one or more future dates and times for a job to execute.	

schedule <recurrence type>

schedule <recurrence type> <interval and date>
no schedule

Sets the type of schedule the job will automatically execute on.
The no form of the command resets the schedule type to its default.

Syntax Description	recurrence type	The available schedule types are: • daily - the job is executed every day at a specified time • weekly - the job is executed on a weekly basis • monthly - the job is executed every month on a specified day of the month • once - the job is executed once at a single specified date and time • periodic - the job is executed on a specified fixed time interval, starting from a fixed point in time.
	interval and date	Interval and date, per recurrence type.
Default	once	
Configuration Mode	config job	
History	3.1.0000	
Role	admin	

Example	`switch (config)# job 100`
	`switch (config job 100) # schedule monthly interval 10`
	`switch (config job 100) #`	

Related Commands	show jobs	
Notes	A schedule type is essentially a structure for specifying one or more future dates and times for a job to execute.	

show jobs

show jobs [<job-id>]

Displays configuration and state (including results of last execution, if any exist) of all jobs, or of one job if a job ID is specified.

Syntax Description	job-id Job ID.
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # show jobs 10`
	`Job 10:`
	`  Status: inactive`
	`  Enabled: yes`
	`  Continue on failure: no`
	`  Schedule Type: once`
	`  Time and date: 1970/01/01 00:00:00 +0000`
	`  Last Exec Time: Thu 2012/04/05 13:11:42 +0000`
	`  Next Exec Time: N/A`
	`  Commands:`
	`  Command 10: show power`
	`  Last Output:`
	`==========`
	`Module Status`
	`==========`
	`PS1 OK`
	`PS2 NOT PRESENT`
	
	`switch (config) #`

Related Commands	show jobs
Notes	

4.16 Statistics and Alarms

4.16.1 Commands

stats alarm <alarm-id> clear

stats alarm <alarm ID> clear

Clears alarm state.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats alarm cpu_util_indiv clear` `switch (config) #`	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> enable

stats alarm <alarm-id> enable
no stats alarm <alarm-id> enable

Enables the alarm.
The no form of the command disables the alarm, notifications will not be received.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
Default	The default is different per alarm-id	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats alarm cpu_util_indiv enable` `switch (config) #`	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> event-repeat

stats alarm <alarm ID> event-repeat {single | while-not-cleared}
no stats alarm <alarm ID> event-repeat

Configures repetition of events from this alarm.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
	single	Does not repeat events: only sends one event whenever the alarm changes state.
	while-not-cleared	Repeats error events until the alarm clears.
Default	single	
Configuration Mode	config	
History	3.1.0000	
Role	monitor/admin	
Example	`switch (config) # stats alarm cpu_util_indiv event-repeat single` `switch (config) #`	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> {rising | falling}

stats alarm <alarm ID> {rising | falling} {clear-threshold | error-threshold} <threshold-value>

Configure alarms thresholds.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
	falling	Configures alarm for when the statistic falls too low.
	rising	Configures alarm for when the statistic rises too high.
	error-threshold	Sets threshold to trigger falling or rising alarm.
	clear-threshold	Sets threshold to clear falling or rising alarm.
	threshold-value	The desired threshold value, different per alarm.
Default	Default is different per alarm-id	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats alarm cpu_util_indiv falling clear-threshold 10` `switch (config) #`	
Related Commands	show stats alarm	
Notes	Not all alarms support all four thresholds.	

stats alarm <alarm-id> rate-limit

stats alarm <alarm ID> rate-limit {count <count-type> <count> | reset | window <window-type> <duration>}

Configures alarms rate limit.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
	count-type	Long medium, or short count (number of alarms).
	reset	Set the count and window durations to default values for this alarm.
	window-type	Long medium, or short count, in seconds.
Default	Short window: 5 alarms in 1 hour Medium window: 20 alarms in 1 day Long window: 50 alarms in 7 days	
Configuration Mode	config	
History	3.1.0000	
Role	monitor/admin	
Example	`switch (config) # stats alarm paging rate-limit window long 2000` `switch (config) #`	
Related Commands	show stats alarm	
Notes		

stats chd <chd-id> clear

stats chd <CHD ID> clear

Clears CHD counters.

Syntax Description	CHD ID	CHD supported by the system, for example: • cpu_util - CPU utilization: percentage of time spent • cpu_util_ave - CPU utilization average: percentage of time spent • cpu_util_day - CPU utilization average: percentage of time spent • disk_device_io_hour - Storage device I/O read/ write statistics for the last hour: bytes • disk_io - Operating system aggregate disk I/O average (KB/sec) • eth_day • eth_hour • eth_ip_day • eth_ip_hour • fs_mnt_day - Filesystem system usage average: bytes • fs_mnt_month - Filesystem system usage average: bytes • fs_mnt_week - Filesystem system usage average: bytes • ib_day • ib_hour • intf_day - Network interface statistics aggregation: bytes • intf_hour - Network interface statistics (same as “interface” sample) • intf_util - Aggregate network utilization across all interfaces • memory_day - Average physical memory usage: bytes • memory_pct - Average physical memory usage • paging - Paging activity: page faults • paging_day - Paging activity: page faults
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats chd memory_day clear` `switch (config) #`	

Related Commands	show stats chd
Notes	

stats chd <chd-id> enable

stats chd <chd-id> enable
no stats chd <chd-id> enable

Enables the CHD.
The no form of the command disables the CHD.

Syntax Description	chd-id	CHD supported by the system, for example: • cpu_util - CPU utilization: percentage of time spent • cpu_util_ave - CPU utilization average: percentage of time spent • cpu_util_day - CPU utilization average: percentage of time spent • disk_device_io_hour - Storage device I/O read/ write statistics for the last hour: bytes • disk_io - Operating system aggregate disk I/O average: KB/sec • eth_day • eth_hour • fs_mnt_day - Filesystem system usage average: bytes • fs_mnt_month - Filesystem system usage average: bytes • fs_mnt_week - Filesystem system usage average: bytes • ib_day • ib_hour • intf_day - Network interface statistics aggregation: bytes • intf_hour - Network interface statistics (same as “interface” sample) • intf_util - Aggregate network utilization across all interfaces • memory_day - Average physical memory usage: bytes • memory_pct - Average physical memory usage • paging - Paging activity: page faults • paging_day - Paging activity: page faults
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Role	monitor/admin	
Example	`switch (config) # stats chd memory_day enable` `switch (config) #`	

Related Commands	show stats chd
Notes	

stats chd <chd-id> compute time

stats chd <CHD ID> compute time {interval | range} <number of seconds>

Sets parameters for when this CHD is computed.

Syntax Description	CHD ID	Possible IDs: • cpu_util - CPU utilization: percentage of time spent • cpu_util_ave - CPU utilization average: percentage of time spent • cpu_util_day - CPU utilization average: percentage of time spent • disk_device_io_hour - Storage device I/O read/ write statistics for the last hour: bytes • disk_io - Operating system aggregate disk I/O average: KB/sec • eth_day • eth_hour • fs_mnt_day - Filesystem system usage average: bytes • fs_mnt_month - Filesystem system usage average: bytes • fs_mnt_week - Filesystem system usage average: bytes • ib_day • ib_hour • intf_day - Network interface statistics aggregation: bytes • intf_hour - Network interface statistics (same as “interface” sample) • intf_util - Aggregate network utilization across all interfaces • memory_day - Average physical memory usage: bytes • memory_pct - Average physical memory usage • paging - Paging activity: page faults • paging_day - Paging activity: page faults
	interval	Specifies calculation interval (how often to do a new calculation) in number of seconds.
	range	Specifies calculation range, in number of seconds.
	number of seconds	Number of seconds.
Default	Different per CHD	
Configuration Mode	config	
History	3.1.0000	
Role	monitor/admin	

Example	`switch (config) # stats chd memory_day compute time interval 120`
	`switch (config) # show stats chd memory_day`
	`CHD "memory_day" (Average physical memory usage: bytes):`
	`Source dataset: sample "memory"`
	`Computation basis: time`
	`Interval: 120 second(s)`
	`Range: 1800 second(s)`
	`switch (config) #`

Related Commands	show stats chd
Notes	

stats sample <sample-id> clear

stats sample <sample ID> clear

Clears sample history.

Syntax Description	sample ID	Possible sample IDs are: • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - Storage device I/O statistics • disk_io - Operating system aggregate disk I/O: KB/sec • eth • eth-abs • eth_ip • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • interface - Network interface statistics • intf_util - Network interface utilization: bytes • memory - System memory utilization: bytes • paging - Paging activity: page faults • power - Power supply usage • power-consumption • temperature - Modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - Mlag counters statistics: counter units • interface-port-channel - Lag counters statistics: counter units
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats sample temperature clear` `switch (config) #`	
Related Commands	show stats sample	
Notes		

stats sample <sample-id> enable

stats sample <sample-id> enable
no states sample <sample-id> enable

Enables the sample.
The no form of the command disables the sample.

Syntax Description	sample-id	Possible sample IDs are: • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - Storage device I/O statistics • disk_io - Operating system aggregate disk I/O: KB/sec • eth • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • interface - Network interface statistics • intf_util - Network interface utilization: bytes • memory - System memory utilization: bytes • paging - Paging activity: page faults • power - Power supply usage • power-consumption • temperature - Modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - Mlag counters statistics: counter units • interface-port-channel - Lag counters statistics: counter units
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # stats sample temperature enable` `switch (config) #`	
Related Commands	show stats sample	
Notes		

stats export

stats export <format> <sample-id>

Exports collected information to a file. Can export extended "interface-ethernet", "interface-port-channel", "interface-mlag-port-channel"& "power" samples.

Syntax Description	sample-id	sample name for which report file should be generated.
	format	format of report file.
Default	N/A	
Configuration Mode	config	
History	3.7.11xx	
Role	Admin	
Example		

```
(config) # stats export csv memory
Generated report file: memory-20001225-105701.csv
(config) # show files stats memory-20001225-105701.csv
#
# Hostname:         r-mgtswd-251
# Report:           Memory utilization
# Time lower bound: none
# Time upper bound: none
# Export time:      2000/12/25 10:57:01 +0000
# System version:   X86_64 antonstu 2018-12-21 09:11:04 x86_64
#
# Statistic group: Physical memory utilization
# Column 0: Timestamp
# Column 1: Free physical memory (kB)
# Column 2: Used physical memory (kB)
# Column 3: Total physical memory (kB)
#
Timestamp,Free,Used,Total
2000/12/25 09:56:54,3378256,3292732,7527320
2000/12/25 09:57:11,3378408,3292640,7527320
2000/12/25 09:57:31,3378104,3292940,7527320
2000/12/25 09:57:54,3376600,3294436,7527320
2000/12/25 09:58:11,3378740,3292280,7527320
2000/12/25 09:58:31,3377840,3293152,7527320
2000/12/25 09:58:51,3377964,3293008,7527320
2000/12/25 09:59:11,3378156,3292800,7527320
...
```

Related Commands	N/A
Notes	

stats sample max-entries

stats sample <sample-id> max-entries [<max-entries>]
[no] stats sample <sample-id> max-entries [<max-entries>]

Sets number of records to be kept in memory for the counter.
No form sets value to default.

Syntax Description	sample-id	sample name for which number of records should be set.
	max-entries	Number of records. Range: 1 - 1000 records. Default of "interface" samples is 100 records
Default	N/A	
Configuration Mode	config	
History	3.7.11xx	
Role	Admin	
Example	`(config) # no stats sample interface-ethernet max-entries` `(config) # stats sample interface-ethernet max-entries 1000`	
Related Commands	N/A	
Notes	Setting a new value will delete all sample history. History does not persist after reboot.	

stats sample interval

stats sample <sample-id> interval [<interval>]
[no] stats sample <sample-id> interval [<interval>]

Sets the sampling interval between taking of sample records.
Added no form. No form sets interval to default value.

Syntax Description	sample-id	sample name for which interval should be set.
	interval	Measured in seconds. Range: 1 - 60*60*24 (24 hours). Default for "interface" samples is 60 seconds.
Default	N/A	
Configuration Mode	config	
History	3.7.11xx	
Role	Admin	
Example	`(config) # no stats sample interface-ethernet interval` `(config) # stats sample interface-ethernet interval 1`	
Related Commands	N/A	
Notes		

stats clear-all

	stats clear all Clears data for all samples, CHDs, and status for all alarms.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # stats clear-all` `switch (config) #`
Related Commands	N/A
Notes	

show stats alarm

show stats alarm [<Alarm ID> [rate-limit]]

Displays status of all alarms or the specified alarm.

Syntax Description	Alarm ID	May be: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
	rate-limit	Displays rate limit parameters.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show stats alarm`
	`Alarm cpu_util_indiv (Average CPU utilization too high):  ok`
	`Alarm disk_io (Operating System Disk I/O per second too high): (disabled)`
	`Alarm fs_mnt (Free filesystem space too low):  ok`
	`Alarm intf_util (Network utilization too high):  (disabled)`
	`Alarm memory_pct_used (Too much memory in use):  (disabled)`
	`Alarm paging (Paging activity too high):  ok`
	`Alarm temperature (Temperature is too high):  ok`
	`switch (config) #`	

Related Commands	stats alarm	
Notes		

show stats chd

show stats chd [<CHD ID>]

Displays configuration of all statistics CHDs.

Syntax Description	CHD ID	May be: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show stats chd disk_device_io_hour`
	
	`CHD "disk_device_io_hour" (Storage device I/O read/write statistics for the last`
	`hour: bytes):`
	`Enabled: yes`
	`Source dataset: sample "disk_device_io"`
	`Computation basis: data points`
	`Interval: 1 data point(s)`
	`Range: 1 data point(s)`
	`switch (config) #`	

Related Commands	stats chd	
Notes		

show stats cpu

	show stats cpu Displays some basic stats about CPU utilization: • the current level • the peak over the past hour • the average over the past hour
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # show stats cpu`
	
	`CPU 0`
	`  Utilization:  6%`
	`  Peak Utilization Last Hour: 16% at 2012/02/28 08:47:32`
	`  Avg. Utilization Last Hour: 8%`
	`switch (config) #`

Related Commands	N/A
Notes	

show stats sample

show stats sample [<sample ID>]

Displays sampling interval for all samples, or the specified one.

Syntax Description	sample ID	Possible sample IDs are: • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - Storage device I/O statistics • disk_io - Operating system aggregate disk I/O: KB/sec • eth • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • interface - Network interface statistics • intf_util - Network interface utilization: bytes • memory - System memory utilization: bytes • paging - Paging activity: page faults • power - Power supply usage • power-consumption • temperature - Modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - Mlag counters statistics: counter units • interface-port-channel - Lag counters statistics: counter units
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show stats sample fan`
	`Sample "fan" (Fan speed):`
	`  Enabled:  yes`
	`  Sampling interval: 1 minute 11 seconds`
	`switch (config) #`	

Related Commands	N/A	
Notes		

show stats sample data

show stats sample <sample-id> data [interface {ethernet | port-channel | mlag-port-channel} <device/port> [counter <counter-name>]] [group name <group-name> [counter <counter-name>]] [max-samples {<max-samples> | all}]

Displays history of counter values. In other words - displays collected information for a sample. Able to limit output information to:

- group of counters
- counter of group of counters

Able to choose number of counter records to be displayed. When there are more records in history than displayed - output for group ends with "…" (ellipses).

Syntax Description	sample ID	sample name for which collected information should be displayed.
	interface/group	Allows limiting output to a particular group of counters. By default, all groups with all counters are displayed. Names of groups are dependent on chosen <sample-id>.
	counter	Allows limiting output to a particular counter. By default all counters related to the chosen group are displayed. This option is available only if "interface/group" option is chosen. Name of particular counter is dependent on <sample-id> and "interface/group".
	max-samples	Allows choosing a number of counter records to display. Range: 1 - 1000 records. The "all" option is meant for all available records. The default is 20 counter records to display.
Default	N/A	
Configuration Mode	N/A	
History	3.7.11xx	
Role	Admin	

Example

```
[standalone: master] > show stats sample interface-ethernet data inter-
face ethernet 1/1 max-samples 1

Sampling data for Interface ethernet counters:
Eth1/1:
    ----------------------------------------------------------------------
    Name                            Timestamp               Value
    ----------------------------------------------------------------------
    Rx_packets                      2000/12/25 10:27:53     0
    Rx_unicast_packets              2000/12/25 10:27:53     0
    Rx_multicast_packets            2000/12/25 10:27:53     0
    Rx_broadcast_packets            2000/12/25 10:27:53     0
    Rx_bytes                        2000/12/25 10:27:53     0
    Rx_discard_packets              2000/12/25 10:27:53     0
    Rx_error_packets                2000/12/25 10:27:53     0
    Rx_fcs_errors                   2000/12/25 10:27:53     0
    Rx_undersize_packets            2000/12/25 10:27:53     0
    Rx_oversize_packets             2000/12/25 10:27:53     0
    Rx_pause_packets                2000/12/25 10:27:53     0
    Rx_unknown_control_opcode       2000/12/25 10:27:53     0
    Rx_symbol_errors                2000/12/25 10:27:53     0
    Rx_packets_of_64_bytes          2000/12/25 10:27:53     0
    Rx_packets_of_65-127_bytes      2000/12/25 10:27:53     0
    Rx_packets_of_128-255_bytes     2000/12/25 10:27:53     0
    Rx_packets_of_256-511_bytes     2000/12/25 10:27:53     0
    Rx_packets_of_512-1023_bytes    2000/12/25 10:27:53     0
    Rx_packets_of_1024-1518_bytes   2000/12/25 10:27:53     0
    Rx_packets_Jumbo                2000/12/25 10:27:53     0
    Tx_packets                      2000/12/25 10:27:53     0
    Tx_unicast_packets              2000/12/25 10:27:53     0
    Tx_multicast_packets            2000/12/25 10:27:53     0
    Tx_broadcast_packets            2000/12/25 10:27:53     0
    Tx_bytes                        2000/12/25 10:27:53     0
    Tx_discard_packets              2000/12/25 10:27:53     0
    Tx_error_packets                2000/12/25 10:27:53     0
    Tx_hoq_discard_packets          2000/12/25 10:27:53     0
    Tx_pause_packets                2000/12/25 10:27:53     0
    Tx_pause_duration               2000/12/25 10:27:53     0
    ...

new JSON:
[standalone: master] >  show stats sample interface-ethernet data
interface ethernet 1/1 max-samples 1 | json-print

[
    {
        "Sampling data for Interface ethernet counters": [
            {
                "Eth1/1": [
                    {
                        "Tx_multicast_packets": [
                            {
                                "Timestamp": "2000/12/25 10:29:53",
                                "Value": "0"
                            }
    ...
old JSON: N/A
```

Related Commands N/A

Notes	• Filtering keyword is dependent on chosen <sample-id>. For convenience, "interface" samples such as "interface-ethernet", "interface-port-channel" & "interface-mlag-port-channel" have interface related keywords for choosing a counters group.
	• Be aware that this is a history of counters. Autocompletion and output can contain information for groups (interfaces) that is not present anymore in the system, and vice versa. If counters are not sampled - they will not appear in the output.
	• Output of collected information is implemented only for the following samples:
	  • interface-port-channel
	  • interface-ethernet
	  • interface-mlag-port-channel
	  • memory
	  • paging
	  • power

4.17 Chassis Management

The chassis manager provides the user access to the following information:

Table 28 - Chassis Manager Information

Accessible Parameters	Description
switch temperatures	Displays system's temperature
power supply voltages	Displays power supplies' voltage levels
fan unit	Displays system fans' status
power unit	Displays system power consumers
Flash memory	Displays information about system memory utilization.

Additionally, it monitors:

- AC power to the PSUs
- DC power out from the PSUs
- Chassis failures

4.17.1 System Health Monitor

The system health monitor scans the system to decide whether or not the system is healthy. When the monitor discovers that one of the system's modules (leaf, spine, fan, or power supply) is in an unhealthy state or returned from an unhealthy state, it notifies the users through the following methods:

- System logs – accessible to the user at any time as they are saved permanently on the system
- Status LEDs – changed by the system health monitor when an error is found in the system and is resolved
- email/SNMP traps – notification on any error found in the system and resolved

4.17.1.1 Re-Notification on Errors

When the system is in an unhealthy state, the system health monitor notifies the user about the current unresolved issue every X seconds. The user can configure the re-notification gap by running the "health notif-cntr <counter>" command.

4.17.1.2 System Health Monitor Alerts Scenarios

- System Health Monitor sends notification alerts in the following cases:

Table 29 - System Health Monitor Alerts Scenarios (Sheet 1 of 2)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
<fan_name> speed is below minimal range	A chassis fan speed is below minimal threshold: 15% of maximum speed	Email, fan LED and system status LED set red, log alert, SNMP.	Check the fan and replace it if required	"<fan_name> has been restored to its normal state"
Fan <fan_number> speed in spine number <spine_number> is below minimal range	A spine fan speed is below minimal threshold: 30% of maximum speed	Email, fan LED and system status LED set red, log alert, SNMP	Check the fan and replace it if required	"Fan speed <fan_number> in spine number <spine_number> has been restored to its normal state"
<fan_name> is unresponsive	A chassis fan is not responsive on Onyx systems	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	"<fan_name> has been restored to its normal state"
Fan <fan_number> in spine number <spine_number> is unresponsive	A spine fan is not responsive on Onyx systems	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	"Fan <fan_number> in spine number <spine_number> has been restored to its normal state"
<fan_name> is not present	A chassis fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	"<fan_name> has been restored to its normal state"
Fan <fan_number> in spine number <spine_number> is not present.	A spine fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	"Fan <fan_number> in spine number <spine_number> has been restored to its normal state"
Insufficient number of working fans in the system	Insufficient number of working fans in the system	Email, fan LED and system status LED set red, log alert, SNMP	Plug in additional fans or change faulty fans	"The system currently has sufficient number of working fans"
Power Supply <ps_number> voltage is out of range	The power supply voltage is out of range.	Email, power supply LED and system status LED set red, log alert, SNMP	Check the power connection of the PS	"Power Supply <ps_number> voltage is in range"

Table 29 - System Health Monitor Alerts Scenarios (Sheet 2 of 2)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
Power supply <ps_number> temperature is too hot	A power supply unit temperature is higher than the maximum threshold of 70 Celsius on Onyx systems	Email, power supply LED and system status LED set red, log alert, SNMP	Check chassis fans connections. On Onyx systems, check system fan connections.	"Power supply <ps_number> temperature is back to normal"
Power Supply <number> is unresponsive	A power supply is malfunctioning or disconnected	Email, system status and power supply LED set red, log alert, SNMP	Connect power cable or replace malfunctioning PS	"Power supply has been removed" or "PS has been restored to its normal state"
Unit/leaf/spine <leaf/spine number> is unresponsive	A leaf/spine is not responsive	Email, system status LED set red, log alert, SNMP	Check leaf/spine connectivity and replace it if required	"Leaf/spine number <leaf/spine number> has been restored to its normal state"
Unit/leaf/spine voltage is out of range	One of the voltages in a Onyx unit is below minimal threshold or higher than the maximum threshold - both thresholds are 15% of the expected voltage	Email, system status LED set red, log alert, SNMP	Check leaf connectivity	"Unit voltage is in range"
ASIC temperature is too hot	A ASIC unit temperature is higher than the maximum threshold of 105 Celsius on Onyx systems.	Email, system status LED set red, log alert, SNMP	Check the fans system	"ASIC temperature is back to normal"

4.17.2 Power Management

4.17.2.1 Width Reduction Power Saving

Link width reduction (LWR) is a Mellanox proprietary power saving feature to be utilized to economize the power usage of the fabric. LWR may be used to manually or automatically configure a certain connection between Mellanox switch systems to lower the width of a link from 4X operation to 1X based on the traffic flow.

LWR is relevant only for 40GbE speeds in which the links are operational at a 4X width.

When "show interfaces" is used, a port's speed appears unchanged even when only one lane is active.

LWR has three operating modes per interface:

- Disabled – LWR does not operate and the link remains in 4X under all circumstances.
- Automatic – the link automatically alternates between 4X and 1X based on traffic flow.
- Force – a port is forced to operate in 1X mode lowering the throughput capability of the port. This mode should be chosen in cases where constant low throughput is expected on the port for a certain time period – after which the port should be configured to one of the other two modes, to allow higher throughput to pass through the port.

See command "power-management width" on page 539.

Table 30 - LWR Configuration Behavior

Switch-A Configuration	Switch-B Configuration	Behavior
Disable	Disable	LWR is disabled.
Disable	Force	Transmission from Switch-B to Switch-A operates at 1X. On the opposite direction, LWR is disabled.
Disable	Auto	Depending on traffic flow, transmission from Switch-B to Switch-A may operate at 1X. On the opposite direction, LWR is disabled.
Auto	Force	Transmission from Switch-B to Switch-A operates at 1 lane. Transmission from Switch-A to Switch-B may operate at 1X depending on the traffic.
Auto	Auto	Width of the connection depends on the traffic flow
Force	Force	Connection between the switches operates at 1x

4.17.3 Monitoring Environmental Conditions

Step 1. Display module's temperature. Run:

```
switch (config) # show temperature
---------------------------------------------------------
Module      Component              Reg  CurTemp    Status
                                        (Celsius)
---------------------------------------------------------
MGMT        SPC                    T1   35.00      OK
MGMT        Board AMB temp         T1   24.00      OK
MGMT        Ports AMB temp         T1   28.00      OK
MGMT        CPU package Sensor     T1   32.00      OK
MGMT        CPU Core Sensor        T1   33.00      OK
MGMT        CPU Core Sensor        T2   30.00      OK
PS1         power-mon              T1   23.00      OK
PS2         power-mon              T1   22.50      OK
```

Step 2. Display measured voltage levels of power supplies. Run:

```
switch (config) # show voltage
-------------------------------------------------------------------------------------
Module   Power Meter          Reg              Expected  Actual   Status  High   Low
                                                Voltage   Voltage          Range  Range
-------------------------------------------------------------------------------------
MGMT     acdc-monitor1        SoC Core         0.97      0.97     OK      1.11   0.82
MGMT     acdc-monitor1        SoC VNN          1.01      1.01     OK      1.16   0.86
MGMT     acdc-monitor1        CPU 0.675V       0.68      0.66     OK      0.78   0.57
MGMT     acdc-monitor1        1V               1.00      0.99     OK      1.15   0.85
MGMT     acdc-monitor1        VDDQ             1.35      1.34     OK      1.55   1.15
MGMT     acdc-monitor1        1.8V             1.80      1.80     OK      2.07   1.53
MGMT     acdc-monitor1        SYS 3.3V         3.30      3.28     OK      3.79   2.80
MGMT     acdc-monitor1        12V              12.00     11.88    OK      13.80  10.20
MGMT     acdc-monitor1        1.35V            1.35      1.34     OK      1.55   1.15
MGMT     acdc-monitor1        VCCSRAM          1.07      1.06     OK      1.23   0.91
MGMT     acdc-monitor1        1.5V             1.50      1.49     OK      1.72   1.27
MGMT     acdc-monitor1        5V               5.00      4.98     OK      5.75   4.25
MGMT     acdc-monitor1        3.3V_AUX         3.30      3.30     OK      3.79   2.80
MGMT     ASICVoltMonitor1     Asic 1.2V        1.20      1.21     OK      1.38   1.02
MGMT     ASICVoltMonitor1     Asic 3.3V        3.30      3.32     OK      3.79   2.80
MGMT     ASICVoltMonitor2     Vcore SPC        0.95      0.96     OK      1.09   0.81
MGMT     ASICVoltMonitor2     Asic 1.8V        1.80      1.81     OK      2.07   1.53
PS1      power-mon            vout 12V         12.00     12.02    OK      13.80  10.20
```

Step 3. Display the fan speed and status. Run:

```
switch (config) # show fan
-------------------------------------------------
Module          Device          Fan  Speed     Status
                                     (RPM)
-------------------------------------------------
FAN1            FAN             F1   9305.00   OK
FAN2            FAN             F1   8823.00   OK
FAN3            FAN             F1   9057.00   OK
FAN4            FAN             F1   9369.00   OK
PS1             FAN             F1   10288.00  OK
PS2             FAN             -    -         NOT PRESENT
```

Step 4. Display the voltage current and status of each module in the system. Run:

```
switch (config) # show power consumers
-------------------------------------------------------------------
Module  Device           Sensor  Power   Voltage  Current  Status
                                 [Watts] [Volts]  [Amp]
-------------------------------------------------------------------
PS1     power-mon        input   37.50   12.02    3.19     OK
MGMT    acdc-monitor2    input   -       -        -        OK

Total power used : 37.50 Watts
```

4.17.4 USB Access

Onyx can access USB devices attached to switch systems. USB devices are automatically recognized and mounted upon insertion. To access a USB device for reading or writing a file, you need to provide the path to the file on the mounted USB device in the following format:

```
scp://username:password@hostname/var/mnt/usb1/<file name>
```

While username and password are the admin username and password and hostname is the IP of the switch.

Examples:

- ***To fetch an image from a USB device, run the command:***

```
switch (config) # “image fetch scp://admin:admin@127.0.0.1/var/mnt/usb1/image.img
```

- ***To save log file ‘my-logfile’ to a USB device under the name*** `test_logfile` ***using the*** `logging files` ***command, run (in Enable or Config mode):***

```
switch (config) # logging files upload my-logfile scp://username:password@hostname/var/
mnt/usb1/test_logfile
```

- ***To safely remove the USB and to flush the cache, after writing (log files, for example) to a USB, use the*** `usb eject` ***command (in Enable or Config mode).***

```
switch (config) # usb eject
```

4.17.5 Unit Identification LED

The unit identification (UID) LED is a hardware feature used as a means of locating a specific switch system in a server room.

- ***To activate the UID LED on a switch system, run:***

```
switch (config) # led MGMT uid on
```

- ***To verify the LED status, run:***

```
switch (config) # show leds
Module   LED           Status
------------------------------------------------------------------------
MGMT     UID           Blue
```

- ***To deactivate the UID LED on a switch system, run:***

```
switch (config) # led MGMT uid off
```

4.17.6 System Reboot

4.17.6.1 Rebooting 1U Switches

- ***To reboot a 1U switch system:***

Step 1. Enter Enable or Config mode. Run:

```
switch >
switch > enable
switch # configure terminal
```

Step 2. Reboot the system. Run:

```
switch (config) # reload
```

4.17.7 Viewing Active Events

Onyx supports viewing all active events on the system. The following events may be observed with the command “show system hardware events”.

Table 31 - Observable System Events

Event Name	Description
Ethernet Family	
Invalid Mac (SMAC=MC)	Source MAC is a multicast address
Invalid Mac (SMAC=DMAC)	Source MAC is same as destination mac address
Invalid Ethertype	Packet has an unknown Ethertype (0x05DC < ethertype < 0x600)
IP Routing Family	
Ingress Router interface is disabled	Ingress packet has been dropped because incoming L3 interface is admin down
Mismatched IP (UC DIP over MC/ BC Mac)	Packet MAC is multicast/broadcast but destination IP is unicast
Invalid IP (DIP=loopback)	Destination IP is loopback IP (For IPv6: DIP==::1/128 or DIP==0:0:0:0:0:ffff:7f00:0/104 For IPv4: DIP==127.0.0.0/8)
Invalid IP (SIP=MC)	Source IP is multicast address (For IPv6: SIP == FF00::/8 For IPv4: SIP == 224.0.0.0: 239.255.255.255 aka 224.0.0.0/4)
Invalid IP (SIP=unspecified)	Source IP is unspecified
Invalid IP (SIP=DIP)	Source IP is identical to destination IP
Mismatched MC Mac	Packet's multicast MAC does not correspond to packet's MC IP address
IPv6 neighbor not resolved	IPv6 neighbor not resolved
Invalid IPv6 (SIP=Link Local)	Source IP is link local (IPv6)
MC RPF check failure	Multicast RPF check failure
TTL expired	TTL value is zero
Egress Router interface is disabled	Egress packet has been dropped because outgoing L3 interface is admin/oper is down
IPv4 neighbor not resolved	Entry not found for destination
Tunnel Family	
NVE Decap fragmentation error	Fragmentation error during decapsulation

4.17.8 Commands

4.17.8.1 Chassis Management

clear counters

clear counters [all | interface <type> <number>] [ethernet | port-channel]

Clears switch counters.

Syntax Description	all	Clears all switch counters
	type	A specific interface type
	number	The interface number
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
	3.6.4000	Added note
Role	admin	
Example	`switch (config) # clear counters`	
Related Commands		
Notes	The command also clears storm-control counters.	

clear system hardware events

	clear system hardware events Clears all active events.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Role	admin
Example	`switch (config) # clear system hardware events`
Related Commands	show system hardware events
Notes	

health

health {max-report-len <length> | re-notif-cntr <counter> | report-clear}

Configures health daemon settings.

Syntax Description	max-report-len <length>	Sets the length of the health report - number of line entries. Range: 10-2048.
	re-notif-cntr <counter>	Health control changes notification counter, in seconds. Range: 120-7200 seconds.
	report-clear	Clears the health report.
Default	max-report-len: 50 re-notif-cntr:	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # health re-notif-cntr 125` `switch (config) #`	
Related Commands	show health-report	
Notes		

led uid

led <module> uid <on | off>

Configures the UID LED.

Syntax Description	module	Specifies the module whose UID LED to configure
	on	Turns on UID LED
	off	Turns off UID LED
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
Role	admin	
Example	`switch (config) # led MGMT uid on` `switch (config) #`	
Related Commands	N/A	
Notes	•	

power-management width

	power-management width {auto | force} **no power-management width** Sets the width of the interface to be automatically adjusted. The no form of the command disables power-saving.	
Syntax Description	auto	Allows the system to automatically decide whether to work in power-saving mode or not.
	force	Forces power-saving mode on the port.
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.3.4000	
Role	admin	
Example	`switch (config  1/1) # power-management width auto`	
Related Commands	show interface	
Notes		

system profile

system profile {eth-default | eth-ipv6-max | eth-ipv4-mc-max} [force]

Optimizes switch system profile to preferred mode.

Syntax Description	eth-default	Balanced Ethernet profile
	eth-ipv6-max	Optimized profile for IPv6 scale
	eth-ipv4-mc-max	Optimized profile for IPv4 multicast scale
	force	Forces operation, without the need for user confirmation
Default	eth-default	
Configuration Mode	config	
History	3.6.6000	
	3.7.11xx	Related command updated
Role	admin	
Example	`switch (config) # system profile eth-default`	
Related Commands	show system profile	
Notes		

usb eject

	usb eject Gracefully turns off the USB interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # usb eject` `switch (config) #`
Related Commands	N/A
Notes	Applicable only for systems with USB interface.

show asic-version

	show asic-version Displays firmware ASIC version.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.2008 Updated Example
Role	admin

Example	`switch (config) # show asic-version`
	`================================================`
	`Module Device Version`
	`================================================`
	`MGMT SPC 15.0200.0092`

Related Commands	N/A
Notes	

show bios

	show bios Displays the BIOS version information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4150
Role	admin

Example	`switch (config) # show bios`
	`BIOS version : 4.6.5`
	`BIOS subversion : Official AMI Release`
	`BIOS release date : 07/02/2013`
	`switch (config) #`

Related Commands	
Notes	

show cpld

	show cpld Displays status of all CPLDs in the system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.4302 Updated example
Role	admin

Example	`switch (config) # show cpld`
	`==================================`
	`Name Type Version`
	`==================================`
	`Cpld1 CPLD_TOR 4`
	`Cpld2 CPLD_PORT1 2`
	`Cpld3 CPLD_PORT2 2`
	`Cpld4 CPLD_MEZZ 3`
	`switch (config) #`

Related Commands	N/A
Notes	

show fan

	show fan Displays fans status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show fan`
	`switch (config) # show fan`
	`==========`
	`Module  Device  Fan  Speed  Status`
	`(RPM)`
	`==========`
	`FAN  FAN  F1  5340.00  OK`
	`FAN  FAN  F2  5340.00  OK`
	`FAN  FAN  F3  5640.00  OK`
	`FAN  FAN  F4  5640.00  OK`
	`PS1  FAN  F1  5730.00  OK`
	`PS2  FAN  -  -  NOT PRESENT`
	`switch (config) #`

Related Commands	N/A
Notes	

show health-report

	show health-report Displays health report.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000 Output update
Role	admin

Example	`switch (config) # show health-report`
	`=======================`
	`| ALERTS CONFIGURATION |`
	`=======================`
	`Re-notification counter (sec):[3600]`
	`Report max counter: [50]`
	`=======================`
	`| HEALTH REPORT |`
	`=======================`
	`No Health issues file`
	`switch (config) #`

Related Commands	N/A
Notes	

show inventory

	show inventory Displays system inventory.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.1604 Removed CPU module output from Example 3.5.1000 Removed Type column from Example 3.6.1002 Updated Example
Role	admin

Example	`switch (config) # show inventory`
	`---------- ---------- ---------- ---------- ----------`
	`Module Part Number Serial Number Asic Rev. HW Rev.`
	`---------- ---------- ---------- ---------- ----------`
	`CHASSIS MSN2100-CB2F MT1752X06330 N/A B3`
	`MGMT MSN2100-CB2F MT1752X06330 1 B3`

Related Commands	N/A
Notes	

show leds

show leds [<module>]

Displays the LED status of the switch system.

Syntax Description	module	Specifies the module whose LED status to display
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.6.2002	Updated Example
Role	admin	
Example		

```
switch (config) # show leds
Module       LED                       Status
-------------------------------------------
MGMT         STATUS                    Green
MGMT         FAN1                      Green
MGMT         FAN2                      Green
MGMT         FAN3                      Green
MGMT         FAN4                      Green
MGMT         PS_STATUS                 Green
MGMT         PS1                       Green
MGMT         PS2                       Green
MGMT         UID                       Blue
```

Related Commands	N/A
Notes	

show memory

	show memory Displays memory status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.1000 Example updated
Role	admin

Example	`scorpion-167 [standalone: master] (config) # show memory`
	
	`----------------------------------------------------------------------`
	`Memory Space  Total       Used        Free        Used+B/C    Free-B/C`
	`----------------------------------------------------------------------`
	`Physical      15848 MB    2849 MB     12999 MB    3854 MB     11994 MB`
	`Swap          0 MB        0 MB        0 MB`
	
	`Physical Memory Borrowed for System Buffers and Cache:`
	`  Buffers            : 27 MB`
	`  Cache              : 910 MB`
	`  Total Buffers/Cache: 937 MB`

Related Commands	N/A
Notes	

show module

	show module Displays modules status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000 Added “Is Fatal” column 3.4.2008 Updated command output 3.4.3000 Updated command output and added note
Role	admin

Example	`switch (config) # show module`
	`=====================`
	` Module    Status`
	`=====================`
	` MGMT      ready`
	` FAN1      ready`
	` FAN2      ready`
	` PS1       ready`
	` PS2       not-present`
	`switch (config) #`

Related Commands	N/A
Notes	The Status column may have one of the following values: error, fatal, not-present, powered-off, powered-on, ready.

show power

	show power Displays power supplies and power usage.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000 Updated Example
Role	admin

Example	`switch (config) # show power`
	`---------------------------------------------------------------------------------`
	`Module Device Sensor Power Voltage Current Capacity Feed Status`
	`[Watts] [Volts] [Amp] [Watts]`
	`---------------------------------------------------------------------------------`
	`PS1 power-mon input 32.25 12.11 1.26 800.00 DC OK`
	`PS2 power-mon input 46.56 12.13 2.33 800.00 DC OK`
	`switch (config) #`

Related Commands	N/A
Notes	

show power consumers

	show power consumers Displays power consumption information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000 Updated Example
Role	admin

Example	`switch (config) # show power consumers`
	`---------------------------------------------------------------------------`
	`Module  Device           Sensor  Power   Voltage  Current  Status`
	`                                  [Watts] [Volts]  [Amp]`
	`---------------------------------------------------------------------------`
	`MGMT    CURR_MONITOR     12V     52.96   11.71    4.52     OK`
	
	`Total power used : 52.96 Watts`
	`switch (config) #`

Related Commands	N/A
Notes	

show protocols

show protocols

Displays all protocols enabled in the system.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.3000
	3.3.4550 Updated Example
	3.6.1002 Updated Example
Role	admin

Example	`switch (config) # show protocols`
	
	`Ethernet enabled`
	`spanning-tree rst`
	`lacp disabled`
	`lldp disabled`
	`igmp-snooping disabled`
	`ets enabled`
	`priority-flow-control disabled`
	`sflow disabled`
	`openflow disabled`
	`mlag disabled`
	`dot1x disabled`
	`isolation-group disabled`
	
	`IP routing disabled`
	`bgp disabled`
	`pim disabled`
	`vrrp disabled`
	`ospf disabled`
	`magp disabled`
	`dhcp-relay disabled`

Related Commands	N/A
Notes	

show resources

	show resources Displays system resources.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show resources`
	`Total  Used  Free`
	`Physical  2027 MB  761 MB  1266 MB`
	`Swap  0 MB  0 MB  0 MB`
	
	`Number of CPUs:  1`
	`CPU load averages: 0.11 / 0.23 / 0.23`
	
	`CPU 1`
	`  Utilization:  5%`
	`  Peak Utilization Last Hour: 19% at 2012/02/15 13:26:19`
	`  Avg. Utilization Last Hour: 7%`
	`switch (config) #`

Related Commands	N/A
Notes	

show system capabilities

	show system capabilities Displays system capabilities.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.3.0000	Added gateway support
	3.6.1002	Updated Example
	3.7.00xx	Updated Example
Role	admin	

Example	`switch (config) # show system capabilities`
	`Ethernet: Supported, L2, L3`
	`Ethernet Max licensed speed: 100Gb`	

Related Commands	show system profile	
Notes		

show system hardware events

show system hardware events <family-name> [clear-on-read]

Displays all active events.

Syntax Description	family-name	Displays all active events per event family: • ethernet • tunnel • ip
	clear-on-read	Clears all active events after displaying them
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
Role	admin	

Example	`switch (config) # show system hardware events clear-on-read`
	
	`Ethernet:  smac is mc;`
	`smac equal dmac;`
	
	`IP:  packet to router is not ip;`
	
	`Tunnel:`	

Related Commands	N/A	
Notes		

show system mac

	show system mac Displays system MAC address.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show system mac`
	`00:02:C9:5E:AF:18`
	`switch (config) #`

Related Commands	N/A
Notes	

show system profile

	show system profile Displays system profile.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0000
Role	admin

Example	`switch (config) # show system profile`
	
	`Profile: eth-default`

Related Commands	system profile
Notes	

show system profile detailed

show system profile detailed

Displays detailed system profile.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000
Role	admin

Example	`switch (config) # show system profile detailed`
	
	`Profile: eth-default`
	
	`----------------------------------------------`
	`Parameter Guaranteed Max Value`
	`----------------------------------------------`
	`FDB size 102400`
	`IPMC-L2 lists 10240`
	`IPMC-L3 lists 10240`
	`IPv4 MC/IGMP routes 10240`
	`IPv4 neighbors 51200`
	`IPv6 neighbors 8192`
	`IPv4 routes 100000`
	`IPv6 shorts 51200`
	`IPv6 routes 21504`
	`VRF 64`
	`RIF 999`

Related Commands	system profile
Notes	

show system type

	show system type Displays system type.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.1000
Role	admin
Example	`switch (config) # show system type` `SN2100`
Related Commands	
Notes	

show temperature

	show temperature Displays system temperature sensors status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show temperature`
	`----------`
	`Module  Component  Reg  CurTemp  Status`
	`(Celsius)`
	`----------`
	`MGMT  SPC  T1  43.00  OK`
	`MGMT  Ports AMB temp  T1  31.00  OK`
	`MGMT  Board AMB temp  T1  30.00  OK`
	`MGMT  CPU Core Sensor  T1  23.00  OK`
	`MGMT  CPU Core Sensor  T2  23.00  OK`
	`MGMT  CPU Core Sensor  T3  24.00  OK`
	`MGMT  CPU Core Sensor  T4  24.00  OK`

Related Commands	N/A
Notes	

show version

	show version Displays version information for the currently running system image.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show version`
	`Product name:  Onyx`
	`Product release:  3.6.8008`
	`Build ID:  #1-dev`
	`Build date:  2018-07-18 13:46:44`
	`Target arch:  x86_64`
	`Target hw:  x86_64`
	`Built by:  jenkins@c5de6027485e`
	`Version summary:  X86_64 3.6.8008 2018-07-18 13:46:44 x86_64`
	
	`Product model:  x86onie`
	`Host ID:  7CFE9058E01E`
	`System UUID:  03000200-0400-0500-0006-000700080009`
	
	`Uptime:  16h 50m 41.260s`
	`CPU load averages:  2.38 / 2.25 / 2.24`
	`Number of CPUs:  2`
	`System memory:  2860 MB used / 12988 MB free / 15848 MB total`
	`Swap:  0 MB used / 0 MB free / 0 MB total`

Related Commands	N/A
Notes	

show version concise

	show version concise Displays concise version information for the currently running system image.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin
Example	`switch (config) # show version concise` `X86_64 3.6.4006 2017-07-03 16:17:39 x86_64`
Related Commands	N/A
Notes	

show voltage

show voltage

Displays voltage level measurements on different sensors.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.5006 Updated Example
Role	admin
Example	
Related Commands	N/A
Notes	

```
switch (config) # show voltage
================================================================================================
Module  Power Meter          Reg                 Expected  Actual   Status  High   Low
                                                 Voltage   Voltage          Range  Range
================================================================================================
MGMT    BOARD_MONITOR        USB 5V sensor       5.00      5.15     OK      5.55   4.45
MGMT    BOARD_MONITOR        Asic I/O sensor     2.27      2.11     OK      2.55   1.99
MGMT    BOARD_MONITOR        1.8V sensor         1.80      1.79     OK      2.03   1.57
MGMT    BOARD_MONITOR        SYS 3.3V sensor     3.30      3.28     OK      3.68   2.92
MGMT    BOARD_MONITOR        CPU 0.9V sensor     0.90      0.93     OK      1.04   0.76
MGMT    BOARD_MONITOR        1.2V sensor         1.20      1.19     OK      1.37   1.03
MGMT    CPU_BOARD_MONITOR    12V sensor          12.00     11.67    OK      13.25  10.75
MGMT    CPU_BOARD_MONITOR    12V sensor          2.50      2.46     OK      2.80   2.20
MGMT    CPU_BOARD_MONITOR    2.5V sensor         3.30      3.26     OK      3.68   2.92
MGMT    CPU_BOARD_MONITOR    SYS 3.3V sensor     3.30      3.24     OK      3.68   2.92
MGMT    CPU_BOARD_MONITOR    SYS 3.3V sensor     1.80      1.79     OK      2.03   1.57
MGMT    CPU_BOARD_MONITOR    1.8V sensor         1.20      1.24     OK      1.37   1.03
switch (config) #
```

show chassis ha

	show chassis ha Displays Chassis HA parameters and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # show chassis ha`
	`2-node HA state:`
	`  Box management IP: 172.30.1.200/16`
	`  interface: mgmt0`
	
	`  local role: master`
	`  local slot: 1`
	`  other state: ready`
	`  reset count: 0`
	`switch (config) #`

Related Commands	chassis ha
Notes	This command is applicable only for director switch systems.

4.18 Network Management Interfaces

4.18.1 SNMP

Simple Network Management Protocol (SNMP), is a network protocol for the management of a network and the monitoring of network devices and their functions. SNMP supports asynchronous event (trap) notifications and queries.

Onyx supports:

- SNMP versions v1, v2c and v3
- SNMP trap notifications
- Standard MIBs
- Mellanox private MIBs

4.18.1.1 Standard MIBs

Table 32 - Standard MIBs – Textual Conventions and Conformance MIBs

MIB	Standard	Comments
INET-ADDRESS-MIB	RFC-4001	
SNMPV2-CONF		
SNMPV2-TC	RFC 2579	
SNMPV2-TM	RFC 3417	
SNMP-USM-AES-MIB	RFC 3826	
IANA-LANGUAGE-MIB	RFC 2591	
IANA-RTPROTO-MIB	RFC 2932	
IANAifType-MIB		
IANA-ADDRESS-FAMILY-NUMBERS-MIB		

Table 33 - Standard MIBs – Chassis and Switch

MIB	Standard	Comments
RFC1213-MIB	RFC 1213	
IF-MIB	RFC 2863	ifXTable only supported.
ENTITY-MIB	RFC 4133	
ENTITY-STATE-MIB	RFC 4268	Fan and temperature states
ENTITY-SENSOR-MIB	RFC 3433	• Port module transmit/receiver power sensors • Fan and temperature sensors
Bridge MIB	RFC 4188	dot1dTpFdbGroup and dot1dStaticGroup are not supported in this MIB, it is supported as a part of Q-Bridge-MIB.

Table 33 - Standard MIBs – Chassis and Switch

MIB	Standard	Comments
Q-Bridge MIB	RFC 4363	The following SNMP groups are not supported: • qBridgeVlanStatisticsGroup, • qBridgeVlanStatisticsOverflowGroup , • qBridgeVlanHCStatisticsGroup, • qBridgeLearningConstraintsGroup. The following SNMP tables are not supported: • dot1qTpGroupTable (dynamic MC MAC addresses) • dot1qForwardAllTable (GMRP) • dot1qForwardUnregisteredTable (GMRP) • dot1qVlanCurrentTable (GVRP)
RSTP-MIB	RFC 4318	
LLDP-MIB	802.1AB-2005	
BGP4-MIB	RFC 4273	Only supports the following tables: • bgpLocalAs • bgpPeerLocalAddr • bgpPeerState • bgpIdentifier
OSPF-MIB	RFC 4750	

4.18.1.2 Private MIB

Table 34 - Private MIBs Supported

MIB	Description
MELLANOX-SMI-MIB	Mellanox Private MIB main structure (no objects)
MELLANOX-PRODUCTS-MIB	List of OID – per managed system (sysObjID)
MELLANOX-IF-VPI-MIB	IfTable extensions
MELLANOX-EFM-MIB	Partially deprecated MIB (based on Mellanox-MIB) Traps definitions and test trap set scalar are supported.
MELLANOX-ENTITY-MIB	Enhances the standard ENTITY-MIB (contains GUID and ASIC revision).
MELLANOX-POWER-CYCLE	Allows rebooting the switch system
MELLANOX-SW-UPDATE-MIB	Allows viewing what SW images are installed, uploading and installing new SW images
MELLANOX-CONFIG-DB	Allows loading, uploading, or deleting configuration files
MELLANOX-ENTITY-STATE-MIB	Extension to support state change traps Note: Currently supported for power supply insertion and extraction only

Table 34 - Private MIBs Supported

MIB	Description
MELLANOX-XSTP-MIB	Extension to support STP information
MELLANOX-DCB-TRAPS	Extension traps for ETC and PFC
MELLANOX-QOS	Proprietary QoS MIBs

Mellanox private MIBs can be downloaded from the Mellanox Support webpage.

4.18.1.3 Proprietary Traps

The following private traps are supported by Onyx.

Table 35 - SNMP MELLANOX-EFM-MIB Traps

Trap	Action Required
asicChipDown	Reboot the system.
asicOverTempReset	Check fans and environmental temperature.
asicOverTemp	Check fans and environmental temperature.
lowPower	Add/connect power supplies.
internalBusError	N/A
procCrash	Generate SysDump and contact Mellanox support.
cpuUtilHigh	N/A
procUnexpectedExit	Generate SysDump and contact Mellanox support.
diskSpaceLow	Clean images and sysDump files using the commands “image delete” and “file debug-dump delete”.
systemHealthStatus	Refer to Health Status table.
lowPowerRecover	N/A
insufficientFans	Check Fans and environmental conditions.
insufficientFansRecover	N/A
insufficientPower	Add/connect power supplies, or change power mode using the command “power redundancy mode”.
insufficientPowerRecover	N/A

For additional information refer to MELLANOX-EFM-MIB.

For event-to-MIB mapping, please refer to Table 26, “Supported Event Notifications and MIB Mapping,” on page 380.

Table 36 - SNMP MELLANOX-POWER-CYCLE Traps

Trap	Action Required
mellanoxPowerCyclePlannedReload	N/A

4.18.1.4 Configuring SNMP

➢ ***To set up the SNMP:***

Step 1. Activate the SNMP server on the Onyx switch (in configure mode) using the following commands:

Community strings are case sensitive.

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
switch (config) # snmp-server community public ro
switch (config) # snmp-server contact "contact name"
switch (config) # snmp-server host <host IP address> traps version 2c public
switch (config) # snmp-server location "location name"
switch (config) # snmp-server user admin v3 enable
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
```

4.18.1.5 Resetting SNMPv3 Engine ID

Resetting SNMP engine ID is not supported on director switch systems.

Switch systems shipped with OS versions older than 3.6.6102 have all had the exact same SNMPv3 engine ID. Going forward, however, all switch systems will ship with a system-specific engine ID.

Upgrading the OS version to 3.6.6102 or higher does not automatically change the current engine ID. That can be done through one of the following methods after performing the software upgrade:

- Running “reset factory”
- Using the command “snmp-server engineID reset” (for more details, please see the procedure below)

➢ ***To reset SNMP engine ID using “snmp-server engineID reset”:***

Prerequisites:

Step 1. If any of the following SNMP configurations exist, please delete/disable them and re-enable/reconfigure them only after SNMP engine ID reset is performed:

1. Make sure SNMP is disabled. Run:

```
switch (config) # no snmp-server enable
```

2. Make sure no SNMP trap host is configured. Run:

```
switch (config) # no snmp-server host <ip-address>
```

3. Make sure no SNMP users are configured. Run:

```
switch (config) # no snmp-server user <username> v3
```

Procedure:

Step 1. Check existing engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <current_key>
```

Step 2. Reset existing engine ID:

```
switch (config) # snmp-server engineID reset
```

Step 3. Verify new engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <new_key>
```

4.18.1.6 Configuring an SNMPv3 User

➢ ***To configure SNMPv3 user:***

Step 1. Configure the user using the command:

```
switch (config) # snmp-server user [role] v3 prompt auth <hash type> priv <privacy
type>
```

where

- user role – admin
- auth type – md5 or sha or sha224 or sha256 or sha384 or sha512
- priv type – des or aes-128 or 3des or aes-192 or aes-256

Step 2. Enter authentication password and its confirmation.

Step 3. Enter privacy password and its confirmation.

```
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
Auth password: ********
      Confirm: ********
Privacy password: ********
         Confirm: ********
switch (config) #
```

To retrieve the system table, run the following SNMP command:

```
snmpwalk -v3 -l authPriv -a MD5 -u admin -A  "<Authentication password>" -x DES -X
"<privacy password>" <system ip> SNMPv2-MIB::system
```

4.18.1.7 Configuring an SNMP Notification

➢ ***To set up the SNMP Notification (traps or informs):***

Step 1. Make sure SNMP and SNMP notification are enable. Run:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
```

Step 2. Configure SNMP host with the desired arguments (IP Address, SNMP version, authentication methods). More than one host can be configured. Each host may have different attributes. Run:

```
switch (config) # snmp-server host 10.134.47.3 traps version 3 user my-username auth
sha my-password
```

Step 3. Verify the SNMP host configuration. Run:

```
switch (config) # show snmp host
Notifications enabled:           yes
Default notification community:  public
Default notification port:       162

Notification sinks:

   10.134.47.3
      Enabled:                yes
      Port:                   162 (default)
      Notification type:      SNMP v3 trap
      Username:               my-username
      Authentication type:    sha
      Privacy type:           aes-128
      Authentication password: (set)
      Privacy password:       (set)
```

Step 4. Configure the desired event to be sent via SNMP. Run:

```
switch (config) # snmp-server notify event interface-up
```

This particular event is used as an example only.

Step 5. Verify the list of traps and informs being sent to out of the system. Run:

```
switch (config) # show snmp events
Events for which traps will be sent:
  asic-chip-down: ASIC (Chip) Down
  cpu-util-high: CPU utilization has risen too high
  disk-space-low: Filesystem free space has fallen too low
  health-module-status: Health module Status
  insufficient-fans: Insufficient amount of fans in system
  insufficient-fans-recover: Insufficient amount of fans in system recovered
  insufficient-power: Insufficient power supply
  interface-down: An interface's link state has changed to down
  interface-up: An interface's link state has changed to up
  internal-bus-error: Internal bus (I2C) Error
  liveness-failure: A process in the system was detected as hung
  low-power: Low power supply
  low-power-recover: Low power supply Recover
  new_root: local bridge became a root bridge
  paging-high: Paging activity has risen too high
  power-redundancy-mismatch: Power redundancy mismatch
  process-crash: A process in the system has crashed
  process-exit: A process in the system unexpectedly exited
  snmp-authtrap: An SNMP v3 request has failed authentication
  topology_change: local bridge triggered a topology change
  unexpected-shutdown: Unexpected system shutdown
```

To print event notifications to the terminal (SSH or CONSOLE) refer to Section 4.7.1, “Monitor,” on page 333.

4.18.1.8 SNMP SET Operations

Onyx allows the user to use SET operations via SNMP interface. This is needed to configure a user/community supporting SET operations.

4.18.1.8.1 Enabling SNMP SET

➢ ***To allow SNMP SET operations using SNMPv1/v2:***

Step 1. Enable SNMP communities. Run:

```
switch (config) # snmp-server enable communities
```

Step 2. Configure a read-write community. Run:

```
switch (config) # snmp-server community my-community-name rw
```

Step 3. Make sure SNMP communities are enabled (they are enabled by default). Make sure "(DISABLED)" does not appear beside "Read-only communities" / "Read-write communities". Run:

```
switch (config) # show snmp

SNMP enabled   : yes
SNMP port      : 161
System contact :
System location:

Read-only communities:
  public

Read-write communities:
  my-community-name

Interface listen enabled: yes

Listen Interfaces:
  Interface: mgmt0

switch (config) # show snmp
No Listen Interfaces.
```

Step 4. Configure this RW community in your MIB browser.

➢ ***To allow SNMP SET operations using SNMPv3:***

Step 1. Create an SNMPv3 user. Run:

```
switch (config) # snmp-server user myuser v3 auth sha <password1> priv aes-128 <pass-
word2>
```

It is possible to use other configuration options not specified in the example above. Please refer to the command "snmp-server user" on page 609 for more information.

Step 2. Make sure the username is enabled for SET access and has admin capability level. Run:

```
switch (config) # show snmp user
User name: myuser
   Enabled overall:         yes
   Authentication type:     sha
   Privacy type:            aes-128
   Authentication password: (set)
   Privacy password:        (set)
   Require privacy:         yes
   SET access:
      Enabled:              yes
      Capability level:     admin
```

Onyx supports the OIDs for SET operation listed in Table 37 which are expanded upon in the following subsections.

Table 37 - Supported SET OIDs

MIB Name	OID Name	OID
MELLANOX-EFM-MIB	sendTestTrapSet	1.3.6.1.4.1.33049.2.1.1.1.6.0
SNMPv2-MIB	sysName	1.3.6.1.2.1.1.5.0
MELLANOX-CONFIG-DB	mellanoxConfigDBCmdExecute mellanoxConfigDBCmdFilename mellanoxConfigDBCmdStatus mellanoxConfigDBCmdStatusString mellanoxConfigDBCmdUri	1.3.6.1.4.1.33049.12.1.1.2.3.0 1.3.6.1.4.1.33049.12.1.1.2.2.0 1.3.6.1.4.1.33049.12.1.1.2.4.0 1.3.6.1.4.1.33049.12.1.1.2.5.0 1.3.6.1.4.1.33049.12.1.1.2.1.0
MELLANOX-POWER-CYCLE	mellanoxPowerCycleCmdExecute mellanoxPowerCycleCmdStatus mellanoxPowerCycleCmdStatusString	1.3.6.1.4.1.33049.10.1.1.2.1.0 1.3.6.1.4.1.33049.10.1.1.2.2.0 1.3.6.1.4.1.33049.10.1.1.2.3.0
MELLANOX-SW-UPDATE	mellanoxSWUpdateCmdSetNext mellanoxSWUpdateCmdUri mellanoxSWUpdateCmdExecute mellanoxSWUpdateCmdStatus mellanoxSWUpdateCmdStatusString mellanoxSWActivePartition mellanoxSWNextBootPartition	1.3.6.1.4.1.33049.11.1.1.2.1.0 1.3.6.1.4.1.33049.11.1.1.2.2.0 1.3.6.1.4.1.33049.11.1.1.2.3.0 1.3.6.1.4.1.33049.11.1.1.2.4.0 1.3.6.1.4.1.33049.11.1.1.2.5.0 1.3.6.1.4.1.33049.11.1.1.3.0.0 1.3.6.1.4.1.33049.11.1.1.4.0.0

4.18.1.8.2Sending a Test Trap SET Request

Onyx allows the user to use test the notification mechanism via SNMP SET. Sending a SET request with the designated OID triggers a test trap.

Prerequisites:

1. Enable SET operations by following the instructions in Section 4.18.1.8.1, “Enabling SNMP SET,” on page 572.
2. Configure host to which to send SNMP notifications.
3. Set a trap receiver in the MIB browser.

➢ ***To send a test trap:***

Step 1. Send a SET request to the switch IP with the OID 1.3.6.1.4.1.33049.2.1.1.1.6.0.

Step 2. Make sure the test trap is received by the aforementioned trap receiver (OID: 1.3.6.1.4.1.33049.2.1.2.13).

4.18.1.8.3Setting Hostname with SNMP

Mellanox supports setting system hostname using an SNMP SET request as described in SNMPv2-MIB (sysName, OID: 1.3.6.1.2.1.1.5.0).

The restrictions on setting a hostname via CLI also apply to setting a hostname through SNMP. Refer to the command “hostname” on page 165 for more information.

4.18.1.8.4Power Cycle with SNMP

supports power cycling its systems using an SNMP SET request as described in MELLANOX-POWER-CYCLE MIB.

Power cycle command is issued via the OID mellanoxPowerCycleCmdExecute. The following options are available:

- Reload – saves any unsaved configuration and reloads the switch
- Reload discard – reboots the system and discards of any unsaved changes
- Reload force – forces an expedited reload on the system even if it is busy without saving unsaved configuration (equals the CLI command `reload force`)
- Reload slave – reloads the slave management on dual management systems (must be executed from the master management module)

On dual management systems it is advised to connect via the BIP to make sure commands are executed from the master management.

4.18.1.8.5Changing Configuration with SNMP

Mellanox supports making configuration changes on its systems using SNMP SET requests. Configuration requests are performed by setting several values (arguments) and then executing a command by setting the value for the relevant operation.

It is possible to set the parameters and execute the commands on the same SNMP request or separate them to several SET operations. Upon executing a command, the values of its arguments remain and can be read using GET commands.

Once a command is executed there may be two types of errors:

- Immediate: This error results in a failure of the SNMP request. This means a critical error in the SNMP request has occurred or that a previous SET request is being executed
- Delayed: The SET request has been accepted by the switch but an error occurred during its execution.

For example, when performing a fetch (download) operation, an immediate error can occur when the given URL is invalid. A delayed error can occur if the download process fails due to network connectivity issues.

The following parameters are arguments are supported:

- Command URI – URI to fetch the configuration file from or upload the file to (for supported URI format please refer to the CLI command “configuration fetch” for more details)
- Config file name – filename to save the configuration file to or to upload to remote location

The following commands are supported:

- BinarySwitchTo – replaces the configuration file with a new binary configuration file. This option fetches the configuration file from the URI provided in the mellanoxConfigD-

BCmdUri and switches to that configuration file. This command should be preceded by a reload command in order for the new configuration to apply.

- TextApply – fetches a configuration file in human-readable format and applies its configuration upon the current configuration.
- BinaryUpload – uploads a binary format configuration file of the current running configuration or an existing configuration file on the switch to the URI in the mellanoxConfigDBCmdUri command. The filename parameter indicates what configuration file on the switch to upload.
- TextUpload – uploads a human-readable configuration file of the current running configuration of an existing configuration file on the switch to the URI in the mellanoxConfigDBCmdUri command. The filename parameter indicates what configuration file on the switch to upload (same as the CLI command `configuration text generate file <filename> upload`).
- ConfigWrite – saves active configuration to a filename on the switch as given in the filename parameter. In case filename is "active", active configuration is saved to the current saved configuration (same as the CLI command `configuration write`).
- BinaryDelete – deletes a binary based configuration file
- TextDelete – deletes a text based configuration file

4.18.1.8.6Upgrading Onyx Software with SNMP

Mellanox supports upgrading Onyx software using an SNMP SET request as described in MELLANOX-SW-UPDATE MIB.

The software upgrade command is issued via the OID mellanoxSWUpdateCmdExecute. The following options are available:

- Update – fetches the image from a specified URI (equivalent to the command "image fetch" followed by "image install")

  The image to update from is defined by the OID mellanoxSWUpdateCmdUri. The restrictions on the URI are identical to what is supported in the CLI command "image fetch" on page 291.
- Set-Next – changes the image for the next boot equivalent to the CLI command "image boot")

  The partition from which to boot is defined by the OID mellanoxSWUpdateCmdSetNext. The parameters for this OID are as follows:
  - 0 – no change
  - 1 – partition 1
  - 2 – partition 2
  - 3 – next partition (default)

Using the OIDs mellanoxSWUpdateCmdStatus and mellanoxSWUpdateCmdStatusString you may view the status of the latest operation performed from the aforementioned in either integer values, or human-readable forms, respectively. The integer values presented may be as follows:

- 0 – no operation
- 1-100 – progress in percentage

- 101 – success
- 200 – failure

4.18.1.9 IF-MIB and Interface Information

Onyx supports displaying information of switch ports, LAG ports, MLAG ports and VLAN interfaces on all systems via SNMP interface. This feature is enabled by default. The interface information is available in the ifTables, ifXTable and mellanoxIfVPITable. Additionally, traps for interface up/down, and internal link suboptimal speed are enabled. The user has the ability to enable one or both of these traps.

Interface up/down traps are sent whenever there is a change in the interface's operational state. These traps are suppressed for internal links when the internal link's speed does not match the configured speed of the link (mismatch condition).

4.18.2 JSON API

JavaScript Object Notation (JSON) is a machine-to-machine data-interchange format which is supported in Onyx CLI.

The JSON API allows executing CLI commands and receiving outputs in JSON format which can be easily parsed by the calling software.

4.18.2.1 Authentication

The JSON API protocol runs over HTTP/HTTPS and uses the existing web authentication mechanism.

In order to access the system via HTTP/HTTPS, an HTTP/HTTPS client is needed to send POST requests to the system.

HTTPS access to the web-based management console needs to be enabled using the command "web https enable" to allow POST requests.

The HTTPS client must first be authenticated by sending a POST request to the following URL:

```
https://<switch-ip-address>/admin/launch?script=rh&template=login&action=login
```

The POST request content should contain the following data:

```
"f_user_id=<user name>&f_password=<user password>"
```

After a successful login, a session id (cookie) is returned to be used for other HTTPS requests in the system.

See Section 4.18.2.6, "JSON Examples," on page 583 for examples.

4.18.2.2 Sending the Request

After successful authentication, the HTTPS client can start sending JSON requests. All requests (POST and GET) should be sent to the following URL:

```
https://<switch-ip-address>/admin/launch?script=json
```

After the request is handled in the system the HTTPS client receives a JSON response with an indication of the request execution result. If there is data resulting from the request, it is returned as part of the response.

See Section 4.18.2.3, “JSON Request Format,” on page 578 for the CLI request format.

See Section 4.18.2.4, “JSON Response Format,” on page 580 for the reply format.

JSON requests may also be sent using the WebUI. For more information on using the WebUI with JSON, please refer to Section 4.18.2.7, “JSON Request Using WebUI,” on page 588.

4.18.2.3 JSON Request Format

4.18.2.3.1JSON Execution Requests

JSON execution requests are HTTPS POST requests that contain CLI commands to be executed in the system.

Execution request can contain a single command or multiple commands to be executed.

Single command execution request format:

```
{
     "cmd": "<CLI command to execute>"
}
```

Example:

```
{
     "cmd": "show  1/1"
}
```

Multiple command execution request format:

```
{
     "commands":["<CLI cmd 1>", "<CLI cmd 2>", … , <CLI cmd n>]
}
```

Example:

```
{
     "commands":
     [
          "show  1/1",
          "show  1/2"
     ]
}
```

In case of a multiple command request, the execution of the commands is done in the order they appear in the execution list. Note that the execution of a multiple command request will be stopped upon first failure. That is, in case the execution of one of the commands fails, none of the remaining commands will be executed.

See Section 4.18.2.6, “JSON Examples,” on page 583 for examples.

Execution Types

Execution requests can be either synchronous (default) or asynchronous.

Synchronous requests will wait for a JSON response from the system. The synchronous request has a defined wait time after which the user will receive a timeout response. The timeout for a synchronous request is configurable by the user and is 30 seconds by default (see the CLI command “json-gw synchronous-request-timeout” on page 619).

Asynchronous requests will return immediately after sending the request with a reply containing a “job_id” key. The user can use the given job ID to later query for request status and execution results. Queries for asynchronous request results are guaranteed to be accessible up to 60 seconds after the request has been completed.

To specify the execution type, the user needs to add the following key to the JSON execution request:

```
"execution_type":"<async|sync>"
```

Example:

```
{
     "execution_type":"async",
     "cmd": "show  1/1"
}
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for examples.

4.18.2.3.2JSON Query Requests

JSON Query requests are HTTPS GET requests that contain a job ID parameter. Using a query request, the user can get information on the current execution state of an ongoing request or the execution results of a completed request. To send a query request, the user should add the following parameters to the JSON URL:

```
job_id=<job number>
```

Example:

```
https://<switch-ip-address>/admin/launch?script=json&job_id=<job number>
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for more examples.

4.18.2.4 JSON Response Format

Set commands normally do not return any data or output. If a set command does return an output, it will be displayed in the “status_message” field.

4.18.2.4.1 Single Command Response Format

The HTTPS POST response format structure is a JSON object consisting of 4 name-value pairs as follows:

```
{
    "executed_command": "<CLI command that was executed>",
    "status" = "<OK|ERROR>",
    "status_message" = "<information on the status received>",
    "data" = {the information that was asked for in the request}
}
```

- executed_command – the CLI command that was executed in the request
- status – the result of the request execution:
  - “OK” if the execution is successful
  - “ERROR” in case of a problem with the execution

  The value type of this key is “string”.
- data – a JSON object containing the information requested. Returns an empty string if there is no data.
- status message – additional information on the received status. May be empty. The value type of this key is “string”.

Example:

```
{
    “executed_command”: “show  1/1
    "status": "OK",
    "status_message": "",
    "data":
    {
        "speed": "40GbE",
        "admin_state": "up"
    }
}
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for more examples.

4.18.2.4.2 Multiple Command Response Format

The HTTPS response format structure is a JSON object consisting of a list of JSON results. Each JSON structure in the list is structured the same as in the single command execution response (see the previous section).

However, the status field can contain in this case an additional value, “ABORTED”, in case a previous command failed. This status value indicates that the command has not been executed at all in the system.

```
{
    "results": [
    {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
        "status_message": "<…>",
        "data": {…}
     },
     {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
        "status_message": "<…>",
        "data": {…}
      },
…
      {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
        "status_message": "<…>",
        "data": {…}
       }
              ]
}
```

Example:

```
{
    "results": [
    {
        "executed_command": "show  1/1",
        "status": "OK",
        "status_message": ""
        "data": {"speed":"40GbE", "admin_state":"up"}
    },
    {
        "executed_command": "show  1/100",
        "status": "ERROR",
        "status_message": "wrong interface name",
        "data": ""
    },
    {
        "executed_command": "show  1/2",
        "status": "ABORTED",
        "status_message": "",
        "data": ""
    }
              ]
}
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for more examples.

4.18.2.4.3Query Response Format

Response to a query request can be of two types. In case the request completes its execution, the response will be similar to the single/multiple command response format, depending on the format of the request, and will display the execution results.

In case the execution is not complete yet, the response format will be similar to the single command response format. However, the status field will contain in this case the value “PENDING” to indicate that the request is still in progress. In addition, the “executed_command” field will contain the current request command being handled by the system.

Example:

```
{
    "executed_command": "show  1/1",
    "status": "PENDING",
    "status_message": "",
    "data":""
}
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for examples.

4.18.2.4.4Asynchronous Response Format

Response to an asynchronous request is similar to the HTTPS response format of the single command response. However, an additional unique field will be added, “job_id”, containing the job id number for querying the request later. The value of the job_id key is of type string.

Another difference is that the “executed_command” field will be empty.

Example:

```
{
    "executed_command": ""
    "status": "OK"
    "status_message": ""
    "data": ""
    "job_id": "2754930426"
}
```

See Section 4.18.2.6, “JSON Examples,” on page 583 for examples.

4.18.2.5 Supported Commands

4.18.2.5.1Set Commands

All non-interactive CLI set commands are supported.

Interactive commands are commands which require user interaction to complete (e.g., type “yes” to confirm). These commands are not supported by the JSON API.

4.18.2.5.2Show Commands

Not all CLI show commands are currently supported by the JSON API. Unsupported commands return an error indication.

Support for all show commands will be completed in future Onyx releases.

For a list “show” commands not currently supported, please refer to Appendix B,“Show Commands Not Supported by JSON,” on page 1943.

4.18.2.6 JSON Examples

The following examples use curl (a common tool in Linux systems) to send HTTPS POST requests to the system.

4.18.2.6.1Authentication Example

Before sending JSON HTTPS request, the user must first authenticate. Run the following from your server’s shell to create a login session ID in the file: /tmp/cookie.

```
curl -c /tmp/cookie -d "f_user_id=admin&f_password=admin"
"https://10.10.10.10/admin/launch?script=rh&template=login&action=login"
```

Upon a successful login, you will receive a reply similar to the following:

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>302 Found</title>
</head><body>
<h1>Found</h1>
<p>The document has moved <a href="https://10.10.10.10/admin/launch?script=rh&amp;tem-
plate=home">here</a>.</p>
<hr>
<address>Apache Server at 10.10.10.10 Port 80</address>
</body></html>
```

The session id can now be used in all other JSON HTTPS requests to the system.

4.18.2.6.2 Synchronous Execution Request Example

Single Command

This example sends a request to query the system profile.

Request (save it to a file named req.json):

```
{"cmd": "show system profile"}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Response:

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
    "status": "OK",
    "executed_command": "show system profile",
    "status_message": "",
    "data": {
        "Profile": "ib",
        "Adaptive Routing": "yes",
        "Number of SWIDs": "1"
    }
}
```

Multiple Commands

This example sends a request to change an interface description and then queries for its status.

Request (save it to a file named req.json):

```
{"commands": ["interface ib 1/1 description test description",
"show interfaces ib 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Response:

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interface ib 1/1 description test description",
            "status_message": "",
            "data": ""
        },
        {
            "status": "OK",
            "executed_command": "show interfaces ib 1/1 status",
            "status_message": "",
            "data": {
                "IB1/1": [
                    {
                        "Description": "test description",
                        "Speed": "fdr",
                        "Logical port state": "Initialize",
                        "Physical port state": "LinkUp",
                        "Current line rate": "56.0 Gbps",
                        "IB Subnet": "infiniband-default"
                    }
                ]
            }
        }
    ]
}
```


4.18.2.6.3Asynchronous Execution Request Example

This example sends an asynchronous request to change an interface description and then queries for its status.

Request (save it to a file named req.json):

```
{"execution_type":"async",
"commands": ["interface ib 1/1 description test description",
"show interfaces ib 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Response:

The system immediately returns a response similar to the following:


```
{
        "executed_command": "",
        "status": "OK",
        "status_message": "",
        "data": "",
        "job_id": "91329386"
}
```

4.18.2.6.4Query Request Example

This example sends a request to query for a job ID received from a previous execution request.

Request:

The request is a an HTTPS GET operation to the JSON URL with the “job_id” parameter.

Send the request:

```
curl -b /tmp/cookie -X GET "https://10.10.10.10/admin/
launch?script=json&job_id=91329386"
```

Response:

If the system is still processing the request, the user receives a response similar to the following:

```
{
        "executed_command": " interface ib 1/1 description test description ",
        "status": "PENDING",
        "status_message": "",
        "data": ""
}
```

If the system is done processing the request, the user receives a response similar to the following:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interface ib 1/1 description test description",
            "status_message": "",
            "data": ""
        },
        {
            "status": "OK",
            "executed_command": "show interfaces ib 1/1 status",
            "status_message": "",
            "data": {
                "IB1/1": [
                    {
                        "Description": "test description",
                        "Speed": "fdr",
                        "Logical port state": "Initialize",
                        "Physical port state": "LinkUp",
                        "Current line rate": "56.0 Gbps",
                        "IB Subnet": "infiniband-default"
                    }
                ]
            }
        }
    ]
}
```

4.18.2.6.5Error Response Example

General Error

This example sends a request with an illegal JSON structure.

Request - without closing bracket “]” (save it to a file named req.json):

```
{"commands": ["interface ib 1/1 description test description",
"show interfaces ib 1/1 status"}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{
    "status": "ERROR",
    "executed_command": "",
    "status_message": "Handle request failed. Reason:\nIllegal JSON structure found in
given JSON data.\nExpecting , delimiter: line 1 column 95 (char 94)",
    "data": ""
}
```

Multiple Command Request Failure

This example sends a multiple command request where one command fails.

Request - with a non-existing interface (1/200) (save it to a file named req.json):

```
{
        "execution_type": "sync",
        "commands": [   "interface ib 1/1 speed sdr",
                        "interface ib 1/200 speed sdr",
                        "interface ib 1/3 speed sdr"]
}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interface ib 1/1 speed sdr",
            "status_message": "",
            "data": ""
        },
        {
            "status": "ERROR",
            "executed_command": "interface ib 1/200 speed sdr",
            "status_message": "% 1st Interface does not exist",
            "data": ""
        },
```

```
        {
            "status": "ABORTED",
            "executed_command": "interface ib 1/3 speed sdr",
            "status_message": "",
            "data": ""
        }
    ]
}
```


4.18.2.7 JSON Request Using WebUI

The Onyx WebUI also allows users to send JSON HTTPS POST and GET requests.

Log into the WebUI, go to the “Setup” tab, and select “JSON API” from the left side menu.

This section is displayed only if JSON API is enabled using the command “json-gw enable”.

4.18.2.7.1To Execute a JSON Request

Step 1. Choose “Execute JSON command”.

Step 2. Choose the “execution_type” from the drop down list.

Step 3. In the “commands” field, type the CLI command(s) to execute.

Use the “+” and “-” buttons to add or remove additional commands to the request.

Step 4. Click “Submit”.

The JSON response is then shown in the “JSON Response” box below.

The HTTPS method (HTTPS POST in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

Figure 12: JSON API WebUI Example

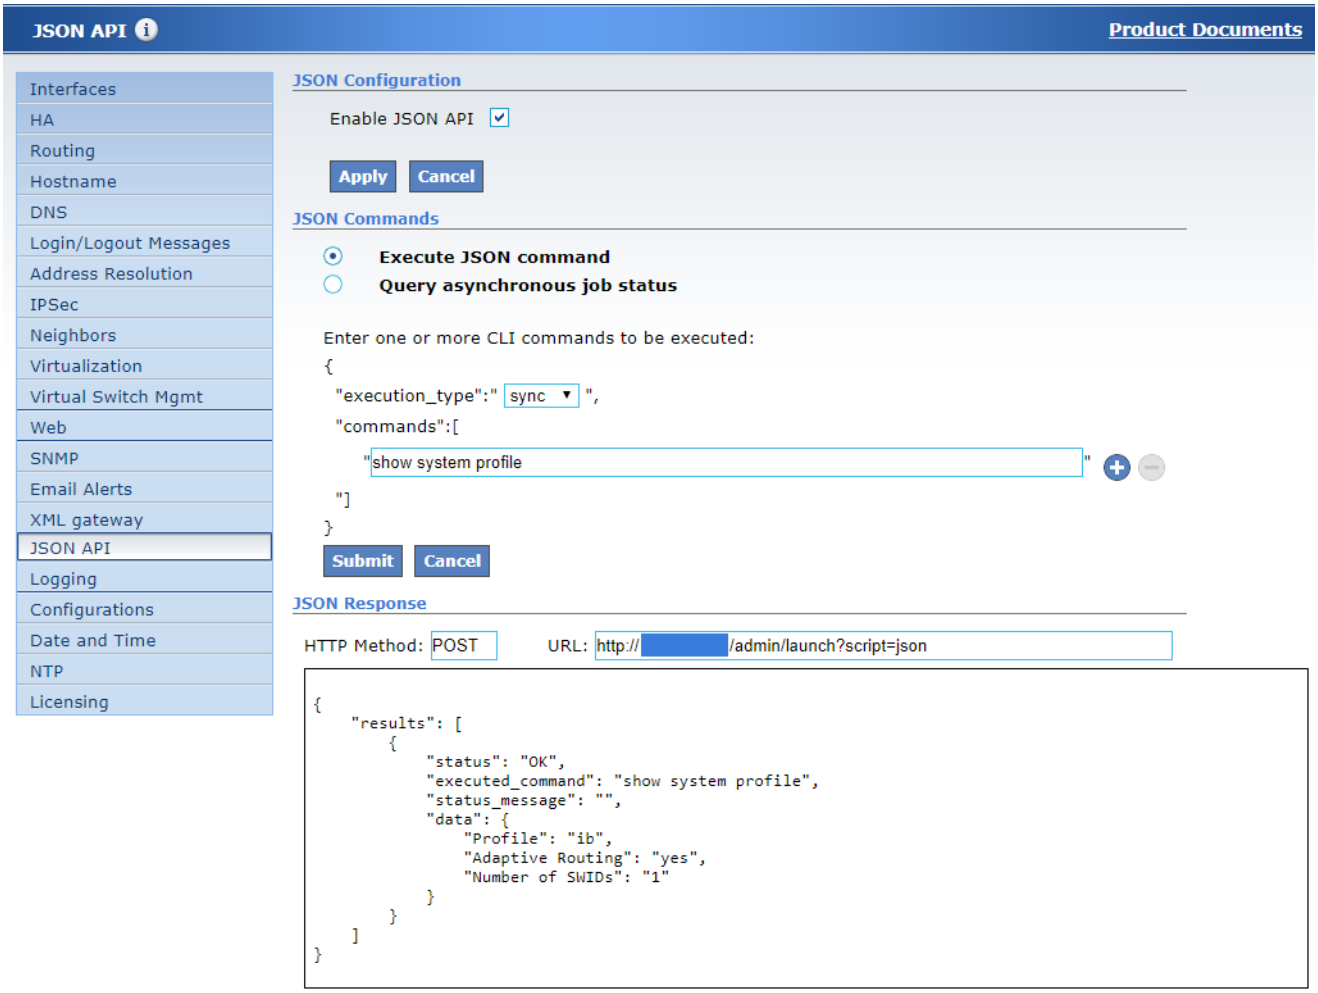

4.18.2.7.2To Query an Asynchronous JSON Request

Step 1. Choose “Query asynchronous job status”.

Step 2. Type the job ID in the “Job ID” text box.

Step 3. Press “Query Status”.

The JSON response is then shown in the “JSON Response” box below.

The HTTPS method (HTTPS GET in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

Figure 13: JSON API Asynchronous Job WebUI Example

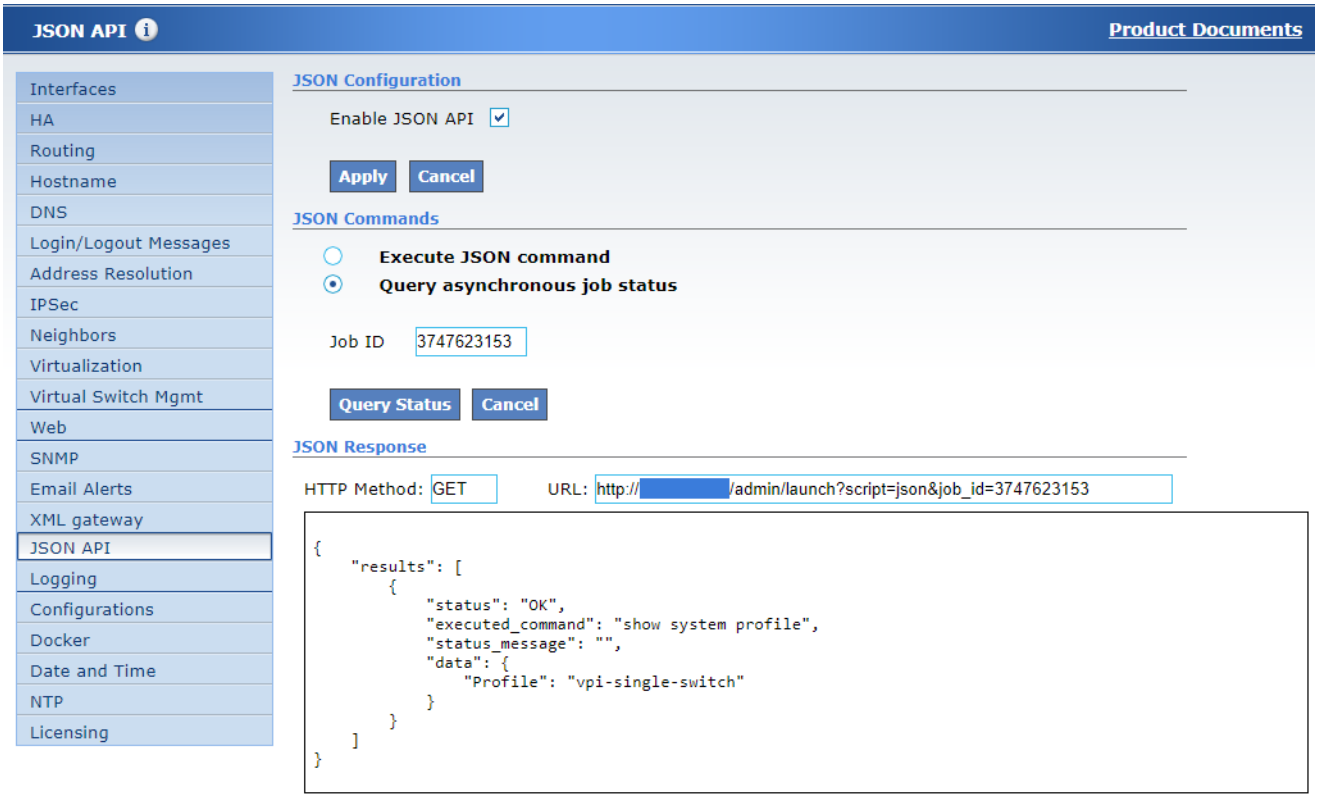

4.18.3 XML API

Onyx XML API is documented in the *Onyx XML API Reference Guide*.

4.18.4 Commands

4.18.4.1 SNMP Commands

The commands in this section are used to manage the SNMP server.

snmp-server auto-refresh

snmp-server auto-refresh {enable | interval <time>}
no snmp-server auto-refresh enable

Configures SNMPD refresh settings.
The no form of the command disables SNMPD refresh mechanism.

Syntax Description	enable	Enables SNMPD refresh mechanism.
	interval	Sets SNMPD refresh interval.
	time	In seconds. Range: 20-500.
Default	Enabled. Interval: 60 secs	
Configuration Mode	config	
History	3.2.3000	
	3.4.1100	Added time parameter and updated notes
Role	admin	
Example	`switch (config) # snmp-server auto-refresh interval 120`	
Related Commands	show snmp	
Notes	• When configuring an interval lower than 60 seconds, the following warning message appears asking for confirmation: “Warning: this configuration may increase CPU utilization, Type 'YES' to confirm: YES”. • When disabling SNMP auto-refresh, information is retrieved no more than once every 60 seconds just like SNMP tables that do not have an auto-refresh mechanism.	

snmp-server cache enable

	[no] snmp-server cache enable Enables/disables snmp cache in case auto-refresh is disabled.
Syntax Description	If snmp cache is disabled, every snmp request will get updated data.
Default	Enabled
Configuration Mode	Configure terminal
History	3.7.00xx
Role	admin
Example	`scorpion2-75 [standalone: master] (config) # snmp-server cache enable`
Related Commands	show snmp auto-refresh [no] snmp-server auto-refresh enable
Notes	If snmp auto-refresh is enabled, the value of cache is meaningless

snmp-server community

snmp-server community <community> [ro | rw]
no snmp-server community <community>

Sets a community name for either read-only or read-write SNMP requests. The no form of the command sets the community string to default.

Syntax Description	community	Community name.
	ro	Sets the read-only community string.
	rw	Sets the read-write community string.
Default	Read-only community: “public” Read-write community: “”	
Configuration Mode	config	
History	3.1.0000	
Role	admin	

Example	`switch(config) # snmp-server community private rw`
	`switch (config) # show snmp`
	`SNMP enabled: yes`
	`SNMP port: 161`
	`System contact:`
	`System location:`
	`Read-only community: public`
	`Read-write community: private`
	
	`Interface listen enabled: yes`
	`No Listen Interfaces.`
	
	`Traps enabled: yes`
	`Default trap community: public`
	`Default trap port: 162`
	
	`No trap sinks configured.`
	`switch(config) #`	

Related Commands	show snmp	
Notes	• If neither the “ro” or the “rw” parameters are specified, the read-only community is set as the default community • If the read-only community is specified, only queries can be performed • If the read-write community is specified, both queries and sets can be performed	

snmp-server contact

snmp-server contact <contact name>
no snmp-server contact

Sets a value for the sysContact variable in MIB-II.
The no form of the command resets the parameter to its default value.

Syntax Description	contact name	Contact name.
Default	“”	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # snmp-server contact my-name`	
Related Commands	show snmp	
Notes		

snmp-server enable

	snmp-server enable **no snmp-server enable** Enables SNMP-related functionality (SNMP engine, and traps) The no form of the command disables the SNMP server.
Syntax Description	N/A
Default	SNMP is enabled by default
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # snmp-server enable`
Related Commands	show snmp
Notes	

snmp-server enable

	snmp-server enable **no snmp-server enable** Enables SNMP-related functionality (SNMP engine, and traps) The no form of the command disables the SNMP server.
Syntax Description	N/A
Default	SNMP is enabled by default
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # snmp-server enable`
Related Commands	show snmp
Notes	

snmp-server engineID reset

	snmp-server engineID reset Resets the SNMPv3 engine ID to be node unique.
Syntax Description	N/A
Default	Default engineID is unchanged
Configuration Mode	config
History	3.6.6102
Role	admin
Example	`switch (config) # snmp-server engienID reset`
Related Commands	show snmp engineID
Notes	Changing system profile or performing “reset factory...” causes the engine ID to change to the new node-unique one.

snmp-server enable mult-communities

	snmp-server enable mult-communities **no snmp-server enable mult-communities** Enables multiple communities to be configured. The no form of the command disables multiple communities to be configured.
Syntax Description	N/A
Default	SNMP server multi-communities are disabled by default
Configuration Mode	config
History	N/A
Role	admin
Example	`switch (config) # snmp-server enable mult-communities`
Related Commands	show snmp
Notes	

snmp-server enable notify

	snmp-server enable notify **no snmp-server enable notify** Enables sending of SNMP traps and informs from this system. The no form of the command disables sending of SNMP traps and informs from this system.
Syntax Description	N/A
Default	SNMP notifies are enabled by default
Configuration Mode	config
History	N/A
Role	admin
Example	`switch (config) # snmp-server enable notify`
Related Commands	show snmp
Notes	SNMP traps are only sent if there are trap sinks configured with the “snmp-server host...” command, and if these trap sinks are themselves enabled.

snmp-server enable set-permission

	snmp-server enable set-permission <MIB-name> **no snmp-server enable set-permission <MIB-name>** Allows SNMP SET requests for items in a specified MIB. The no form of the command disallows SNMP SET requests for items in a specified MIB.
Syntax Description	N/A
Default	SNMP MIBs are all given permission for SET requests by default
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # snmp-server enable set-permission MELLANOX-SW-UPDATE`
Related Commands	show snmp set-permission
Notes	

snmp-server host disable

	snmp-server host <ip-address> disable **no snmp-server host <ip-address> [disable]** Temporarily disables sending of all notifications to this host. The no form of the commands resumes sending of all notifications to this host
Syntax Description	IP address IPv4 or IPv6 address
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # snmp-server host 10.10.10.10 disable`
Related Commands	show snmp snmp-server enable
Notes	

snmp-server host informs

snmp-server host <ip-address> informs [<community> | port <port> | version 2c | version 3 {engineID <engineID> | user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]] | encrypted auth ... | prompt auth ...}}]
no snmp-server host <ip-address> informs port

Send SNMP v2c informs to this host with the default trap community.
The no form of the commands removes a host from which SNMP traps should be sent.

Syntax Description	IP address	IPv4 or IPv6 address
	community	Specifies trap community string
	port	Overrides default UDP port for this trap sink
	version	Specifies the SNMP version of traps to send to this host
	engineID	Specifies engine ID of this inform sink
	user	Specifies username for this inform sink
	auth	Configures SNMP v3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted)
	hash-type	• MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user
	privacy-type	• aes-128 – uses AES-128 encryption for privacy • des – uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.
	encrypted	Configure SNMP v3 security parameters, specifying passwords in encrypted form
	prompt	Configure SNMP v3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line
Default	Default community is “public” Default UDP port is 162 Default SNMP version is 2	

Configuration Mode	config
History	3.2.1050
Role	admin
Example	`switch (config) # snmp-server host 1.1.1.1 informs version 3 engineID 0x800041da04643265363932653432303135 user test auth md5 password priv aes-128 password`
Related Commands	show snmp snmp-server enable snmp-server host informs version 3
Notes	

snmp-server host traps

snmp-server host <ip-address> traps [<community> | port <port> | version {1 | 2c} | version 3 {user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]] | encrypted auth ... | prompt auth ...}}]
no snmp-server host <ip-address> traps port

Send SNMP v2c traps to this host with the default trap community.
The no form of the commands removes a host from which SNMP traps should be sent.

Syntax Description	IP address	IPv4 or IPv6 address
	community	Specifies trap community string
	port	Overrides default UDP port for this trap sink
	version	Specifies the SNMP version of traps to send to this host
	user	Specifies username for this inform sink
	auth	Configures SNMP v3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted)
	hash-type	• MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user
	privacy-type	• aes-128 – uses AES-128 encryption for privacy • des – uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.
	encrypted	Configure SNMP v3 security parameters, specifying passwords in encrypted form
	prompt	Configure SNMP v3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line
Default	Default community is “public” Default UDP port is 162 Default SNMP version is 2	
Configuration Mode	config	

History	3.1.0000
Role	admin
Example	`switch (config) # snmp-server host 1.1.1.1 informs version 3 user test auth md5 password priv aes-128 password`
Related Commands	show snmp snmp-server enable snmp-server host informs version 3
Notes	

snmp-server listen

snmp-server listen {enable | interface <ifName>}
no snmp-server listen {enable | interface <ifName>}

Configures SNMP server interface access restrictions.
The no form of the command disables the listen interface restricted list for SNMP server.

Syntax Description	enable	Enables SNMP interface restrictions on access to this system.
	ifName	Adds an interface to the “listen” list for SNMP server. For example: “mgmt0”, “mgmt1”.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # snmp listen enable`	
Related Commands	show snmp	
Notes	If enabled, and if at least one of the interfaces listed is eligible to be a listen interface, then SNMP requests will only be accepted on those interfaces. Otherwise, SNMP requests are accepted on any interface.	

snmp-server notify

snmp-server notify {community <community> | event <event name> | port <port> | send-test}
no snmp-server notify {community | event <event name> | port}

Configures SNMP notifications (traps and informs).
The no form of the commands negate the SNMP notifications.

Syntax Description	community	Sets the default community for traps sent to hosts which do not have a custom community string set.
	event	Specifies which events will be sent as traps.
	port	Sets the default port to which traps are sent.
	send-test	Sends a test trap.
Default	Community: public All informs and traps are enabled Port: 162	
Configuration Mode	config	
History	3.1.0000	
	3.2.1050	Changed traps to notify
Role	admin	
Example	`switch (config) # snmp-server community public`	
Related Commands	show snmp show snmp events	
Notes	• This setting is only meaningful if traps are enabled, though the list of hosts may still be edited if traps are disabled • Refer to Mellanox MIB file for the list of supported traps	

snmp-server port

snmp-server port <port>
no snmp-server port

Sets the UDP listening port for the SNMP agent.
The no form of the command resets the parameter to its default value.

Syntax Description	port	UDP port.
Default	161	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # snmp-server port 1000`	
Related Commands	show snmp	
Notes		

snmp-server user

snmp-server user {admin | <username>} v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] | capability <cap> | enable <sets> | prompt auth <hash-type> [priv <privacy-type>] | require-privacy}
no snmp-server user {admin | <username> } v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] | capability <cap> | enable <sets> | prompt auth <hash-type> [priv <privacy-type>]}

Specifies an existing username, or a new one to be added.
The no form of the command disables access via SNMP v3 for the specified user.

Syntax Description	v3	Configures SNMP v3 users
	auth	Configures SNMP v3 security parameters, specifying passwords in plaintext on the command line (note: passwords are always stored encrypted). Available hash-type options are: <md5|sha|sha224|sha256|sha384|sha512>.
	capability	Sets capability level for SET requests
	enable	Enables SNMP v3 access for this user
	encrypted	Configures SNMP v3 security parameters, specifying passwords in encrypted form
	prompt	Configures SNMP v3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line
	require-privacy	Requires privacy (encryption) for requests from this user
	priv	Configures SNMP v3 security parameters, specifying which protocol to use for traffic encryption. Available priv-type options: <des|3des|aes-128|aes-192|aes-256>.
Default	No SNMP v3 users defined	
Configuration Mode	config	
History	3.1.0000	
	3.7.00xx	
Role	admin	
Example	`switch (config) # snmp-server user admin v3 enable`	

Related Commands	show snmp user
Notes	• The username chosen here may be anything that is valid as a local UNIX username (alphanumeric, plus '-', '_', and '.'), but these usernames are unrelated to, and independent of, local user accounts. That is, they need not have the same capability level as a local user account of the same name. Note that these usernames should not be longer than 31 characters, or they will not work. • The hash algorithm specified is used both to create digests of the authentication and privacy passwords for storage in configuration, and also in HMAC form for the authentication protocol itself. • There are three variants of the command, which branch out after the “v3” keyword. If “auth” is used next, the passwords are specified in plaintext on the command line. If “encrypted” is used next, the passwords are specified encrypted (hashed) on the command line. If “prompt-pass” is used, the passwords are not specified on the command line the user is prompted for them when the command is executing. If “priv” is not specified, only the auth password is prompted for. If “priv” is specified, the privacy password is prompted for; entering an empty string for this prompt will result in using the same password specified for authentication. • AES privacy type encryption using the newest algorithm, which means we use aes-blumenthal. For more information see - http://www.snmp.com/eso/esoConsortiumMIB.txt • No more than 30 SNMP V3 users are allowed in the database.

show snmp

show snmp [events | host]

Displays SNMP-server configuration and status.

Syntax Description	events	SNMP events
	host	List of notification sinks
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show snmp`
	
	`SNMP enabled   : no`
	`SNMP port      : 161`
	`System contact : Test`
	`System location: Boston`
	
	`Read-only communities:`
	`  public`
	
	`Read-write communities:`
	`  good`
	
	`Interface listen enabled: yes`
	
	`Listen Interfaces:`
	`  Interface: mgmt0`	

Related Commands	show snmp	
Notes		

show snmp auto-refresh

	show snmp auto-refresh Displays SNMPD refresh mechanism status.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.6000 Updated Example 3.7.00xx Updated Example
Role	admin

Example	`switch (config) # show snmp auto-refresh`
	`SNMP auto refresh:`
	`Auto-refresh enabled: yes`
	`Refresh interval (sec): 60`
	`Cache enabled: yes`
	
	`Auto-Refreshed tables:`
	`ifTable`
	`ifXTable`
	`mellanoxIfVPITable`

Related Commands	snmp-server auto-refresh
Notes	

show snmp engineID

	show snmp engineID Displays SNMPv3 engine ID key.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6102
Role	admin
Example	`switch (config) # show snmp engineID` `Local SNMP engineID: 0x80004f4db1dd435e80accf4a4d4d3031`
Related Commands	snmp-server engineID
Notes	

show snmp set-permission

	show snmp set-permission Displays SNMP SET permission settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin

Example	`switch (config) # show snmp set-permission`
	`---------------------------------------------`
	`MIB Name                          Set Enable`
	`---------------------------------------------`
	`MELLANOX-CONFIG-DB-MIB            yes`
	`MELLANOX-EFM-MIB                 yes`
	`MELLANOX-POWER-CYCLE             yes`
	`MELLANOX-SW-UPDATE               no`
	`RFC1213-MIB                       no`

Related Commands	snmp-server enable set-permission
Notes	

show snmp user

	show snmp user Displays SNMP user information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.8008 Updated Example
Role	admin

Example	`switch (config) # show snmp user`
	`User name: Hendrix`
	`   Enabled overall:         yes`
	`   Authentication type:     sha`
	`   Privacy type:            des`
	`   Authentication password: (set)`
	`   Privacy password:        (set)`
	`   Require privacy: yes`
	`   SET access:`
	`      Enabled:              yes`
	`      Capability level:     admin`
	`switch (config) #`

Related Commands	show snmp
Notes	

4.18.4.2 XML API Commands

xml-gw enable

	xml-gw enable **no xml-gw enable** Enables the XML gateway. The no form of the command disables the XML gateway.
Syntax Description	N/A
Default	XML Gateway is enabled
Configuration Mode	config
History	3.1.0000
Role	admin

Example	`switch (config) # xml-gw enable`
	`switch (config) # show xml-gw`
	`XML Gateway enabled:  yes`
	`switch (config) #`

Related Commands	show xml-gw
Notes	

show xml-gw

	show xml-gw Displays the XML gateway setting.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Role	admin

Example	`switch (config) # show xml-gw`
	`XML Gateway enabled:  yes`
	`switch (config) #`

Related Commands	xml-gw enable
Notes	

4.18.4.3 JSON API Commands

json-gw enable

	json-gw enable **no json-gw enable** Enables the JSON API. The no form of the command disables the JSON API.
Syntax Description	N/A
Default	JSON API is enabled
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # json-gw enable`
Related Commands	show json-gw
Notes	

json-gw synchronous-request-timeout

	json-gw synchronous-request-timeout <time out value> **no json-gw synchronous-request-timeout** Defines a timeout value for synchronous JSON requests (in seconds). The no form of the command returns the timeout value to its default.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004 3.6.4000 Updated Example and Related Commands.
Role	admin

Example	`switch (config) # show json-gw`
	
	`JSON Gateway enabled:  yes`
	`Synchronous request timeout:  30`
	`JSON API version:  1.0`

Related Commands	json-gw enable json-gw synchronous-request-timeout <time out value> no json-gw synchronous-request-timeout
Notes	

show json-gw

	show json-gw Displays the JSON API setting.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004 3.6.4000 Updated Example and Related Commands.
Role	admin

Example	`switch (config) # show json-gw`
	
	`JSON Gateway enabled:  yes`
	`Synchronous request timeout:  30`
	`JSON API version:  1.0`

Related Commands	json-gw enable json-gw synchronous-request-timeout <time out value> no json-gw synchronous-request-timeout
Notes	

4.19 Puppet Agent

Puppet is a software that allows network administrators to automate repetitive tasks. Onyx includes a built-in agent for the open-source "Puppet" configuration change management system. The Puppet agent enables configuring Mellanox switches in accordance with the standard "puppet-netdev-stdlib" type library and with the "Mellanox-netdev-stdlib-mlnxos" and "Mellanox-netdev-ospf-stdlib" type libraries provided by Mellanox Technologies to the Puppet community.

For more information, please refer to the CLI commands, to the NetDev documentation at https://github.com/puppetlabs/puppet-netdev-stdlib and to Mellanox's Puppet modules GitHub page at https://github.com/Mellanox.

4.19.1 Setting the Puppet Server

➢ ***To set the puppet server:***

Step 1. Define the Puppet server (the name has to be a DNS and not IP). Run:

```
switch (config) # puppet-agent master-hostname <please_type_your_hostname_DNS_here>
switch (config) #
```

Step 2. Enable the Puppet agent. Run:

```
switch (config) # puppet-agent enable
switch (config) #
```

Step 3. (Optional) Verify there are no errors in the Puppet agent log. Run:

```
switch (config) # show puppet-agent log continuous
switch (config) #
```

4.19.2 Accepting the Switch Request

This is to be performed on the first run only.

➢ ***To accept the switch's request:***

Option 1 – using Puppet CLI commands:

Step 1. Ensure the certificate request. Run:

```
# puppet cert list
"<switch>"
(F4:B4:20:3B:2B:11:76:37:14:34:D0:D1:03:ED:3D:B5)
```

Step 2. Sign the certificate request if the cert_name parameter (e.g. switch1.domain) is in the list. Run:

```
# puppet cert sign <full_domain_name>
```

Step 3. Verify the request is removed from the Puppet certification list. Run:

```
# puppet cert list
```

Option 2 – accept certificate requests in the puppet server console:

Step 1. Go to the "nodes requests" page (the button is at the top right), and wait for a certificate request for the switch and then accept it.

Figure 14: Accepting an Agent Request through the Console

4.19.3 Installing Modules on the Puppet Server

Mellanox uses netdev-stdlib types and provides a package of Mellanox providers for those types which have to be installed at the Puppet server prior to the first Puppet configuration run (before configuring resources on the Mellanox switch).

To install those modules, run the following commands in the Puppet server:

```
# puppet module install netdevops-netdev_stdlib
# puppet module install mellanox-netdev_ospf_stdlib
# puppet module install mellanox-netdev_stdlib_mlnxos
```

In case of an already installed module, please use the command "`puppet module upgrade <module_name>`" or "`puppet module install <module_name> --force`" instead of "`puppet module install <module_name>`" to reinstall the modules.

For more information please refer to the Network Automation Tools document or Puppet category in the Mellanox community site at: http://community.mellanox.com/community/support/solutions.

4.19.4 Writing Configuration Classes

➢ ***To write configuration classes:***

Step 1. Assigning Configuration Classes to a Node

Configuration files can be written and changed in the puppet server machine in the directory "/etc/puppetlabs/puppet/manifests/" (or "/etc/puppet/manifests" in case of an open source puppet server).

The file “/etc/puppetlabs/puppet/manifests/site.pp” is the main file for Puppet-classes-to-nodes association. To associate a configuration to a Puppet agent node, just append association lines as below:

```
import "netdev_vlan_example"
import "netdev_l2_vlan_example"
import "netdev_lag_example"
node 'switch-6375dc.mtr.labs.mlnx'{

  netdev_device { $hostname: }

  include vlan_example # Asserts a class vlan_example in one of the files
  include l2_interface_example

 include lag_example

}
```

If you have a puppet console, you may assign classes of configuration in the following way:

- Add the relevant classes (using the console add class button on the “nodes” page).
- Assign the classes to the relevant nodes/groups in the puppet server console (in the console node/group page -> edit -> Classes).

Step 2. Update VLAN.

Manifest example (located in “/etc/puppetlabs/puppet/manifests/netdev_vlan_example.pp”).

```
class vlan_example{

  $vlans = {
   'Vlan244' => {vlan_id => 244, ensure => present},
   'Vlan245' => {vlan_id => 245, ensure => present},
  }

  create_resources( netdev_vlan, $vlans )
}
```

Step 3. Update Layer 2 Interface.

Manifest example (located in "/etc/puppetlabs/puppet/manifests/netdev_l2_interface_example.pp")

```
class vlans_ensure_example{

  $vlans = {
   'Vlan347' => {vlan_id => 347, ensure => present},
   'Vlan348' => {vlan_id => 348, ensure => present},
   'Vlan349' => {vlan_id => 349, ensure => present},
  }

  create_resources( netdev_vlan, $vlans )
}

class l2_interface_example{

  include vlans_ensure_example #class to Ensure VLANs before assigning

  $l2_interfaces = {
    'ethernet 1/3' => {ensure => absent, vlan_tagging => disable}, #default
    'ethernet 1/4' => {ensure => present, vlan_tagging => enable,
tagged_vlans => [Vlan348,Vlan347], untagged_vlan => Vlan349} #hybrid
  }

  create_resources( netdev_l2_interface, $l2_interfaces )
}
```

Step 4. Update LAG.

Manifest example (located in "/etc/puppetlabs/puppet/manifests/netdev_lag_example.pp")

```
class lag_example{

  $lags = {
  'port-channel 101' => {ensure => present,
links => ['ethernet 1/12', 'ethernet 1/13'], lacp => active},
  'port-channel 102' => {ensure => present,
links => ['ethernet 1/6','ethernet 1/5'], lacp => disabled},
  }

  create_resources( netdev_lag, $lags )
}
```

You may add classes to ensure that all assigned links are with the same layer 1 and layer 2 configurations (similarly to the way we did in update l2_interface section with vlans_ensure_example class).

4.19.5 Supported Configuration Capabilities

4.19.5.1 Ethernet and Port-Channel Interface Capabilities

Table 38 - Ethernet and Port-Channel Interface Capabilities

Field	Description	Values	Example
ensure	Sets the given values or restores the interface to default	absent, present	ensure => present
speed	Sets the speed of the interface.	auto*|10m|100m|1g|10g|40g|56g	speed => 1g
admin	Disables/enables interface admin state.	up, down	admin => up
mtu	Configures the maximum transmission unit frame size for the interface.	1518-9216	mtu => 1520
description	Sets the Ethernet and LAG description.	Text	description => “changed_by_puppet”

4.19.5.2 VLAN Capabilities

Table 39 - VLAN Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the VLAN given as a resource ID	absent, present	ensure => present
vlan_id	The VLAN ID	1-4094 (integer)	vlan_id => 245

4.19.5.3 Layer 2 Ethernet Interface Capabilities

Table 40 - L2 Ethernet and Port-Channel Interface Capabilities

Field	Description	Values	Example
ensure	Sets the given values or restores the Layer 2 interface to default.	absent, present	ensure => present
vlan_tagging	VLAN tagging mode	enable,disable	vlan_tagging => enable
tagged_vlans	List of tagged (trunked) VLANs	2-4994 (range)	tagged_vlans => [Vlan348,Vlan347]
untagged_vlan	Untag (access) VLAN	<VLAN name>	untagged_vlan => Vlan349

4.19.5.4 LAG (Port-Channel) Capabilities

Table 41 - LAG Capabilities

Field	Description	Values	Example
ensure	creates or destroys the port-channel given as a resource ID	absent, present	ensure => present
lacp	The LACP mode of the LAG	passive | active | on	lacp => on
links	List of ports assigned to the LAG	List of link names	links => ['ethernet 1/6','ethernet 1/5']

4.19.5.5 Layer 3 Interface Capabilities

Table 42 - L3 Interface Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the interface VLAN specified in the resource ID.	present, absent	ensure => present
ipaddress	Sets IP address on the Layer 3 interface (requires netmask).	A valid IP address	ipaddress => '192.168.4.2'
netmask	Sets netmask for the IP address.	A valid netmask (of the form X.1X2.X3.X4), which creates a valid combination with the given IP address	netmask => '255.255.255.0'
method	Configures the method of the L3 interface (currently supports only static method).	static	method => static

4.19.5.6 OSPF Interface Capabilities

Table 43 - OSPF Interface Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the OSPF interface of the associated interface of the VLAN specified in the resource ID	present, absent	ensure => present
area_id	The associated area ID	Integer representing an IP	area_id => '7200'
Type	The network type	broadcast, point_to_point	type => 'point_to_point'

4.19.5.7 OSPF Area Capabilities

Table 44 - OSPF Area Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the OSPF area specified in the resource ID	present, absent	ensure => present
router_id	The OSPF area associated router ID (currently supports only default router)	default	router_id => 'default'
ospf_area_mode	The OSPF area mode	normal, stub, nssa	ospf_area_mode => 'stub'
subnets	A list of associated subnets	List of subnets	["192.168.4.0/24", "192.168.5.0/24"]

4.19.5.8 Router OSPF Capabilities

Table 45 - Router OSPF Capabilities

Field	Description	Values	Example
ensure	Enables/disables the router ID specified in the resource ID	present, absent	ensure => present

4.19.5.9 SNMP, LLDP, IP Routing, and Spanning Tree Capabilities

Table 46 - Protocol Enable/Disable Capabilities

Field	Description	Values	Example
ensure	Enables/disables the protocol specified in the resource ID	present, absent	ensure => present

4.19.5.10Fetched Image Capabilities

Table 47 - Fetched Image Capabilities

Field	Description	Values	Example
ensure	Enables/disables the protocol specified in the resource ID	present, absent	ensure => present
protocol	Specifies the protocol for fetch method	http, https, ftp, tftp, scp, sftp	protocol => scp
host	The host where the filename located	DNS/IP	host => my_DNS

Table 47 - Fetched Image Capabilities

Field	Description	Values	Example
user	The username for fetching the image	Username	user => my_username
password	The password for fetching the image	Password	password => my_pass
location	The location of the file name in the host file system	Directory full path	location => '/tmp'
force_delete	Remove all the images or only the ones which are not installed on any partition, before fetching	yes, no	force_delete => no

4.19.5.11 Installed Image Capabilities

Table 48 - Installed Image Capabilities

Field	Description	Values	Example
ensure	Specifies if the image version given in as resource ID is ensured to be installed or not	present, absent	ensure => present
is_next_boot	Ensures that the installed image is the next boot partition	yes, no	is_next_boot => yes
configuration_write	Writes configurations to database.	yes, no	configuration_write => yes
force_reload	Reload if image is in other partition.	yes, no	force_reload => no

4.19.6 Supported Resources for Each Type

Table 49 - Fetched Image Capabilities

Resource Type	Puppet Type Name	Supported Resource IDS	Example
Network device	netdev_device	$hostname	netdev_device { $hostname: }
Layer 1 interface	netdev_interface	'ethernet <#ID>', 'port-channel <#id>', 'ib <#ID>'	netdev_interface{'ethernet 1/3': ensure => absent}
Layer 2 interface	netdev_l2_interface	'ethernet <#ID>', 'port-channel <#id>'	netdev_l2_interface{'ethernet 1/3': ensure => absent}

Table 49 - Fetched Image Capabilities

Resource Type	Puppet Type Name	Supported Resource IDS	Example
VLAN	netdev_vlan	VLAN name string	netdev_vlan {'Vlan244': vlan_id => 244, ensure => present }
LAG	netdev_lag	'port-channel <#id>'	netdev_lag {'port-channel 101': ensure => present }
Layer 3 interface	netdev_l3_interface	'vlan <#ID>'	netdev_l3_interface{ 'vlan 4': ipaddress => '192.168.4.2', netmask => '255.255.255.0'}
OSPF interface	netdev_ospf_interface	'vlan <#ID>'	netdev_ospf _interface{ 'vlan 4': ensure => present, area_id => '10' }
OSPF area	netdev_ospf_area	Valid area ID (representing an IP)	netdev_ospf_area{'10': ensure => present, ospf_area_mode=>'stub'}
OSPF router	netdev_router_ospf	Currently only supports 'default'	netdev_router_ospf {'default': ensure => present }
Protocol	mlnx_protocol	ip_routing, lldp, snmp, spanning_tree	mlnx_protocol { 'ip_routing': ensure => present}
Fetched image	mlnx_fetched_img	The image file name	mlnx_fetched_image { 'onyx-X86_64-3.6.8008.img': ensure => present}
Installed image	mlnx_installed_img	The image version name	mlnx_installed_img { '3.3.4300': ensure => present}

4.19.7 Troubleshooting

This section presents common issues that may prevent the switch from connecting to the puppet server.

4.19.7.1 Switch and Server Clocks are not Synchronized

This can be fixed by using NTP to synchronize the clocks at the switch (using the CLI command `ntp`) and at the server (e.g. using `ntpdate`).

4.19.7.2 Outdated or Invalid SSL Certificates Either on the Switch or the Server

This can be fixed on the switch using the CLI command `puppet-agent clear-certificates` (requires `puppet-agent restart` to take effect).

On the server it can be fixed by running `puppet cert clean <switch_fqdn>` (FQDN is the Fully Qualified Domain Name which consists of a hostname and a domain suffix).

4.19.7.3 Communications Issue

Make sure it is possible to ping the puppet server hostname from the switch (using the CLI command `ping`).

If the hostname is not reachable (e.g. no DNS server) it can be statically added to the switch local hosts lookup (using the CLI command `ip host`).

Make sure that port 8140 is open (using the command tracepath {<hostname> | <ip>}/8140).

4.19.8 Commands

puppet-agent

	puppet-agent Enters puppet agent configuration mode.
Syntax Description	N/A
Default	None
Configuration Mode	config
History	3.3.4200
Role	admin
Example	`switch (config) # puppet-agent` `switch (config puppet-agent) #`
Related Commands	
Notes	

master-hostname

master-hostname <hostname>
no master-hostname

Sets the puppet server hostname.
The no form of the command resets the parameter to its default.

Syntax Description	hostname	Puppet server hostname. Free string may be entered.
Default	puppet	
Configuration Mode	config puppet	
History	3.3.4200	
Role	admin	
Example	`switch (config puppet-agent) # master-hostname my-puppet-server-host-name` `switch (config puppet-agent) #`	
Related Commands		
Notes		

enable

	enable **no enable** Enables the puppet server on the switch. The no form of the command disables the puppet server.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config puppet
History	3.3.4200
Role	admin
Example	`switch (config puppet-agent) # enable` `switch (config puppet-agent) #`
Related Commands	
Notes	

run-interval

run-interval <time>

Configures the time interval in which the puppet agent reports to the puppet server.

Syntax Description	time	Can be in seconds (“30” or “30s”), minutes (“30m”), hours (“6h”), days (“2d”), or years (“5y”).
Default	30m	
Configuration Mode	config puppet	
History	3.3.4302	
Role	admin	
Example	`switch (config puppet-agent) # run-interval 40m` `switch (config puppet-agent) #`	
Related Commands	show puppet-agent	
Notes		

restart

	puppet-agent restart Restarts the puppet agent.
Syntax Description	N/A
Default	N/A
Configuration Mode	config puppet
History	3.3.4200
Role	admin
Example	`switch (config puppet-agent) # restart` `switch (config puppet-agent) #`
Related Commands	
Notes	

show puppet-agent

	show puppet-agent Displays Puppet agent status and configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4200 3.3.4302 Updated Example with “Run interval” 3.7.00xx Updated Example with “Puppet agent: disabled”
Role	admin

Example	`switch (config puppet-agent) # show puppet-agent`
	`Puppet agent: disabled`
	`Puppet master hostname: puppet`
	`Run interval: 30m`
	`switch (config puppet-agent) #`

Related Commands	
Notes	

show puppet-agent log

show puppet-agent log [[not] [matching | continuous] <string> | files [[not] matching] <string>]

Displays the Puppet agent's log file.

Syntax Description	continuous	Puppet agent log messages as they arrive.
	files	Displays archived Puppet agent log files.
	matching	Displays Puppet agent log that match a given string.
	not	Displays Puppet agent log that do not meet a certain string.
	string	Free string.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4200	
Role	admin	
Example		
Related Commands		
Notes		

```
switch (config puppet-agent) # show puppet-agent log
Mon Nov 04 11:52:42 +0000 2013 Puppet (notice): Starting Puppet client version 3.2.3
Mon Nov 04 11:52:44 +0000 2013 Puppet (warning): Unable to fetch my node definition, but the agent run will continue:
Mon Nov 04 11:52:44 +0000 2013 Puppet (warning): Could not intern from pson: source '"#<Puppet::Node:0x7f' not in PSON!
Mon Nov 04 11:53:21 +0000 2013 /Netdev_vlan[Vlan104]/ensure (notice): created
Mon Nov 04 11:53:22 +0000 2013 /Netdev_vlan[Vlan101]/ensure (notice): created
Mon Nov 04 11:53:23 +0000 2013 /Netdev_vlan[Vlan102]/ensure (notice): created
Mon Nov 04 11:53:24 +0000 2013 /Netdev_vlan[Vlan103]/ensure (notice): created
Mon Nov 04 11:53:40 +0000 2013 /Netdev_l2_interface[ethernet 1/6]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan103'
Mon Nov 04 11:53:43 +0000 2013 /Netdev_l2_interface[ethernet 1/7]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan103'
Mon Nov 04 11:53:48 +0000 2013 /Netdev_vlan[Vlan100]/ensure (notice): created
Mon Nov 04 11:53:48 +0000 2013 /Netdev_l2_interface[ethernet 1/5]/vlan_tagging (notice): vlan_tagging changed 'enable' to 'disable'
Mon Nov 04 11:53:48 +0000 2013 /Netdev_l2_interface[ethernet 1/5]/tagged_vlans (notice): tagged_vlans changed '[]' to
'[Vlan100,Vlan101,Vlan102]'
Mon Nov 04 11:53:51 +0000 2013 /Netdev_l2_interface[ethernet 1/1]/tagged_vlans (notice): tagged_vlans changed '[]' to '[Vlan101,Vlan104]'
Mon Nov 04 11:53:51 +0000 2013 /Netdev_l2_interface[ethernet 1/1]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100'
Mon Nov 04 11:53:54 +0000 2013 /Netdev_l2_interface[ethernet 1/3]/tagged_vlans (notice): tagged_vlans changed '[]' to '[Vlan101,Vlan104]'
Mon Nov 04 11:53:54 +0000 2013 /Netdev_l2_interface[ethernet 1/3]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100'
Mon Nov 04 11:53:58 +0000 2013 /Netdev_l2_interface[ethernet 1/4]/vlan_tagging (notice): vlan_tagging changed 'enable' to 'disable'
Mon Nov 04 11:53:58 +0000 2013 /Netdev_l2_interface[ethernet 1/4]/tagged_vlans (notice): tagged_vlans changed '[]' to
'[Vlan100,Vlan101,Vlan102]'
Mon Nov 04 11:54:03 +0000 2013 /Netdev_l2_interface[ethernet 1/2]/tagged_vlans (notice): tagged_vlans changed '[]' to '[Vlan101,Vlan104]'
Mon Nov 04 11:54:03 +0000 2013 /Netdev_l2_interface[ethernet 1/2]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100'
Mon Nov 04 11:54:06 +0000 2013 Puppet (notice): Finished catalog run in 47.90 seconds
switch (config puppet-agent) #
```

4.20 Virtual Machine

A virtual machine (VM) on a switch is added to allow additional OS to run on top of the switch. The VM OS can connect through mgmt0 interface to the switch system's management interface. In addition, the VM is also connected to the out-of-band network. This allows it to communicate through the network and to control the switch management software.

The number of VMs that may run on a system is user-configurable and also relies on resource availability.

The number of configurable VMs is limited to 4.

Each VM consumes the following resources:

- Memory
- Processing power which is not policed (the user may determine the core to be used)
- MACs which are required for each vNIC (user configurable)

4.20.1 Virtual Machine Configuration

➢ ***To configure a VM:***

The example below installs Ubuntu 14 and defines 3GB storage with 512MB memory (default) using the first core of the switch system (default) through mgmt0 interface (default) with an auto-generated MAC (default).

Step 1. Enable the VM feature. Run:

```
switch (config) # virtual-machine enable
```

Step 2. Create a VM. Run:

```
switch (config) # virtual-machine host my-vm
switch (config virtual-machine host my-vm) #
```

Step 3. Define storage for the VM. Run:

```
switch (config virtual-machine host my-vm) # storage create disk size-max 3000
100.0% [##################################################################]
Created empty virtual disk volume 'vdisk001.img' in pool 'default'
Device attached to drive number 1.
switch (config virtual-machine host my-vm) #
```

Step 4. Display the VM parameters (notice boldface). Run:

```
switch (config virtual-machine host my-vm) # show virtual-machine host my-vm
VM 'my-vm'
   Status:       shut off                  Architecture:     x86_64
   VCPU used:    0 sec                     Number of VCPUs:  1
   Boot order:   hd, cdrom                 Memory size:      512 MB
   Consoles:     text, graphics
   Storage:
      IDE bus, drive 1: default/vdisk001.img (3000 MB capacity)
   Interfaces:
      1: on bridge 'mgmt0'            address unknown   (MAC 52:54:00:2F:89:69)
switch (config virtual-machine host my-vm) # exit
switch (config) #
```

Step 5. Import the VM image. Run:

```
switch (config) # virtual-machine volume fetch url scp://root@<ip>/.../ubuntu-14.04-
server-amd64.iso
Password (if required): *************
 100.0%  [#########################################################################]
```

Step 6. Install the imported image. Run:

```
switch (config) # virtual-machine host my-vm
switch (config virtual-machine host my-vm) # install cdrom file ubuntu-14.04-server-
amd64.iso
```

Step 7. Switch to a different terminal, and run the following command to connect VNC viewer to the VM:

```
$ vncviewer -via admin@<switch IP> 127.0.0.1:0
...
Mellanox Onyx Switch Management

Password: ************
```

Continue VM installation from the VNC prompt.

The switch prompt is unresponsive pending a successful VM installation. Successful VM installation is indicated by the reboot of the VM.

VM IP is determined by DHCP configuration according to the MAC address in Step 4.

➢ ***To verify VM configuration, run:***

```
switch (config virtual-machine host my-vm) # show virtual-machine host my-vm
VM 'my-vm'
   Status:       running                   Architecture:     x86_64
   VCPU used:    12 min 27.440 sec         Number of VCPUs:  1
   Boot order:   cdrom, hd                 Memory size:      512 MB
   Consoles:     text, graphics
   Storage:
      IDE bus, drive 1: default/vdisk001.img (3000 MB capacity)
      IDE bus, drive 2: default/ubuntu-14.04-server-amd64.iso (564 MB capacity) READ-
ONLY
   Interfaces:
      1: on bridge 'mgmt0'           address unknown   (MAC 52:54:00:2F:89:69)
```

➢ ***To perform a VM installation from a USB stick:***

USB stick with supported VM image should be supplied to the user by Mellanox.

Step 1. Insert the USB stick (supplied by Mellanox) to the USB port of your switch system.

Step 2. Decide on a name for the VM (e.g. “my_vm”).

Step 3. Decide on the network configuration of the VM.

- Use DHCP or alternately use static IP definitions
- Assign a MAC address or alternately use the default MAC address

Step 4. Launch the full installation of the VM with the network definitions of your choice.

4.20.2 Commands

4.20.2.1 Config

virtual-machine enable

virtual-machine enable
no virtual-machine enable

Enables VM feature on the switch.
The no form of the command disables VM feature on the switch.

Syntax Description	N/A
Default	no virtual-machine enable
Configuration Mode	config
History	3.4.0000
Role	admin
Example	`switch (config) # virtual-machine enable`
Related Commands	
Notes	

virtual-machine host

virtual-machine host <vm-name>
no virtual-machine host <vm-name>

Creates a VM, or enters its configuration context if it already exists.
The no form of the command removes the VM with the specified name.

Syntax Description	vm-name	Configures a name for the VM
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# virtual-machine host my-vm` `switch (config virtual-machine host my-vm)#`	
Related Commands		
Notes		

arch

arch {i386 | x86_64}

Configures VM CPU architecture.

Syntax Description	i386	32-bit x86 CPU architecture
	x86_64	64-bit x86 CPU architecture
Default	x86_64	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# arch i386`	
Related Commands	virtual-machine	
Notes		

comment

comment <string>
no comment

Configures a comment describing the VM.
The no form of the command deletes the configured comment.

Syntax Description	string	Free string
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# comment “example VM”`	
Related Commands	virtual-machine	
Notes	To configure a multi-word string, the string must be placed within quotation marks.	

console

console {connect [graphics | text [force]] | graphics vnc | text tty}
no console {graphics vnc | text tty}

Configures or connects to a text or graphical console.
The no form of the command clears console settings.

Syntax Description	connect	Connects to the text console unless specified otherwise: • graphics – connects to the X11 graphical (VNC) console • text – connects to the text console
	graphics vnc	Enables graphical (VNC) console access
	text tty	Enables TTY text console access
Default	Graphical and textual consoles are enabled	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# console connect text`	
Related Commands	virtual-machine ssh server x11-forwarding enable	
Notes	• To exit the text console press Ctrl-6 (or Ctrl-Shift-6) • If the guest OS is not configured to receive input from a serial console (ttyS0), the VM console becomes unresponsive when connected to. • To view the graphical console, X display must be enabled. There are two options to activate it, the command `vncviewer -via admin@<switchIP> 127.0.0.1:<VNC display num>` (which is run from an external Linux host) and the command `ssh server x11-forwarding enable` (which is run from within the switch and requires that you log out and log back in again using `ssh -X`). The latter command weakens the switch security, therefore, it is recommended to opt for the second option. The VNC display num parameter may be procured by running the command `show virtual-machine <vm-name> detail`.	

install

install {cancel |cdrom [pool <pool-name>] {file <volume-name> [connect-console <console-type> | disk-overwrite | timeout {<minutes> | none}]}}

Installs an operating system onto this VM (temporarily attach a CD and boot from it).

Syntax Description	cancel	Cancels an install already in progress
	cdrom	Installs an operating system from a CD-ROM (ISO) image
	pool <pool-name>	Configures storage pool in which to find image to install: • default • usb
	file <volume-name>	Specifies CD-ROM (ISO) image from which to install
	connect-console <console-type>	Connects to the console during installation. The types may be: • text – text console • graphics – graphical console
	disk-overwrite	Installs even if primary target volume is not empty
	timeout {<minutes> | none}	Configures a timeout for installation in minutes (default is no timeout).
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# install cdrom pool usb file <image>`	
Related Commands	virtual-machine	
Notes	The default pool from which the system installs the ISO image is the /var/ partition in the switch.	

install-from-usb

install-from-usb [ip-address <ip-address> <mask> default-gateway <gw-ip> [mac <mac-address>] | mac <mac-address>]

Installs a VM including resource allocation and network configurations from a VM image file located on a USB stick.

Syntax Description	ip-address	The IP address to configure for the installed VM
	mask	The IP mask to configure to the installed VM Format example: /24 or 255.255.255.0 Note that a space is required between the IP address and the netmask length
	default-gateway	The IP address of the default gateway to configure for the installed VM
	mac	The MAC address to configure for the installed VM (e.g. ff:ee:dd:cc:bb:aa)
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.6.2002	
Role	admin	

Example	`switch (config virtual-machine host my-vm)# install-from-usb`
	`100.0% [##################################################################]`
	`VM host my-vm MAC is: aa:bb:cc:dd:ee:ff`
	`switch (config virtual-machine host my-vm)#`	

Related Commands	virtual-machine	
Notes	USB stick supplied by Mellanox must be inserted into the USB port of the switch system prior to running this command.	

interface

interface <id> {bridge <bridge> | macaddr <mac> | model <model> | name <name>}

Configures virtual interfaces.

Syntax Description	<id>	Interface ID number (1-8 permitted)
	bridge <bridge>	Configures bridge for this interface (i.e. mgmt0 or mgmt1)
	macaddr <mac>	Configures MAC address (e.g. ff:ee:dd:cc:bb:aa)
	model <model>	Configures virtual interface model: • realtek-8139 – Realtek 8139 (default) • virtio – Virtual IO
	name <name>	Configures virtual interface name. The name must begin with “vif”.
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# interface 1 model virtio`	
Related Commands	virtual-machine	
Notes		

memory

memory <MB>

Configures memory allowance.

Syntax Description	MB	Size in megabytes.
Default	512MB	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# memory 1024`	
Related Commands	virtual-machine	
Notes	It is recommended not to allocate more than 1GB of memory per VM.	

power

power {cycle [force | connect-console {graphics | text}] | off [force] | on [connect-console {graphics | text}]}

Turns the VM on or off, or other related options.

Syntax Description	cycle	Powers the VM down and then on again immediately
	force	Forces an action on the system.
	connect-console <console-type>	Connects to the console after power-on. The types may be: • text – text console • graphics – graphical console
	off	Powers down the VM
	on	Powers on VM:
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# power cycle force`	
Related Commands	virtual-machine	
Notes		

storage create

storage create disk [drive-number <number> | file <filename> | mode {read-only | read-write} | pool <pool-name> | size-max <MB>]

Creates a new storage device for the VM, with an automatically assigned name.

Syntax Description	create disk	Creates a new virtual disk image for this VM.
	drive-number <number>	Specifies the drive number to be assigned to the volume. Insert “new” to assign a new drive number to the volume.
	file <filename>	Specifies filename for new volume to be created
	mode {read-only | read-write}	Specifies initial device mode
	pool <pool-name>	Specifies storage pool in which to create new volume
	size-max <MB>	Specifies maximum disk capacity in megabytes
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# storage create disk size-max 2000`	
Related Commands	virtual-machine	
Notes		

storage device

storage device [bus ide] drive-number <number> [mode {read-only | read-write}] source {[pool <pool-name>] file <filename>}
no storage device [bus ide] drive-number <id>

Modifies existing storage device, or create a new one with a specific name.
The no form of the command removes a storage device from the VM.

Syntax Description	device	Modifies existing storage device, or creates a new one with a specific name
	bus ide	Configures bus type to IDE
	drive-number <number>	Selects device to configure by drive number
	mode {read-only | read-write}	Configures the device mode: • read-only – sets the read-only attribute of the volume • read-write – sets the read-write attribute of the volume
	source	Specifies where the data for this volume resides
	file <filename>	Specifies the filename for this volume
	pool <pool-name> file <filename>	Specifies the storage pool for this volume
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# storage create disk bus ide`	
Related Commands	virtual-machine	
Notes		

vcpus

vcpus {count <count> | vcpu <vcpu> pin <cpu-list> [<cpu-list>]}
no vcpus {pin | vcpu <vcpu> pin}

Specifies virtual CPUs.
The no form of the command removes certain CPU configuration.

Syntax Description	count <count>	Specifies the number of virtual CPUs
	vcpu <vcpu>	Specifies options for a particular virtual CPU
	pin <cpu-list>	Specifies physical CPUs to pin to this vCPU
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config virtual-machine host my-vm)# vcpus count 1`	
Related Commands		
Notes		

virtual-machine volume fetch url

virt volume fetch url <download-url> [filename <filename> | pool <pool-name> filename <filename>]

Fetches volume image from a remote host.

Syntax Description	download-url	Specifies URL from which to fetch a volume. Format: http, https, ftp, tftp, scp and sftp are supported (e.g. scp://username[:password]@hostname/path/filename)
	filename <filename>	Specifies new filename for fetched volume image
	pool-name <pool-name>	Specifies storage pool for fetched volume image
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config) # virtual-machine volume fetch scp://admin[:admin-pass]@<hostname/path/filename>`	
Related Commands		
Notes		

virt volume file

virt volume file <name> {create disk size-max <MB> | move {new-name <new-name> | pool <pool-name> new-name <new-name>} | upload <upload-url>}
no virt volume file <volume-name>

Specifies name of volume file to manage.
The no form of the command deletes the volume file.

Syntax Description	file <name>	Specifies name of volume file to manage
	create	Creates a new volume file under this name
	disk size-max <MB>	Specifies maximum capacity of virtual disk to create
	move	Moves or renames this volume
	new-name <filename>	Specifies a name for the destination file
	pool <pool-name> new-name <filename>	Specifies a storage pool for the copy
	upload <upload-url>	Uploads this volume file to a remote host. Format: ftp, tftp, scp and sftp are supported (e.g. scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config virtual machine host	
History	3.4.0000	
Role	admin	
Example	`switch (config) # virt volume file my-vm_file create cdrom extract cdrom1`	
Related Commands		
Notes		

4.20.2.2 Show

show virtual-machine configured

	show virtual-machine configured Displays global virtualization configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.0000
Role	admin

Example	`switch (config) # show virtual-machine configured`
	`Virtualization enabled:  yes`
	`Virtual machines:  2 configured`
	`Virtual networks:  0 configured`
	`switch (config) #`

Related Commands	
Notes	

show virtual-machine host

show virtual-machine host [<vm-name>]

Displays status for this VM.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Role	admin	

Example	`switch (config) # show virtual-machine host my-vm`
	`VM 'my-vm'`
	`Status: shut off Architecture: x86_64`
	`VCPU used: 0 sec Number of VCPUs: 1`
	`Boot order: hd, cdrom Memory size: 512 MB`
	`Consoles: text, graphics`
	`Storage:`
	`IDE bus, drive 1: default/vdisk001.img (3000 MB capacity)`
	`Interfaces:`
	`1: on bridge 'mgmt0' address unknown (MAC 52:54:00:2F:89:69)`
	`switch (config) #`	

Related Commands		
Notes	If the command is run in the middle of an installation, the following banner appears: `*** INSTALL IN PROGRESS: begun <time> ago ***`	

show virtual-machine host configured

show virtual-machine host <vm-name> configured [detail]

Displays configuration for this VM.

Syntax Description	vm-name	The name of the VM.
	detail	Displays detailed configuration for this VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Role	admin	
Example		

```
switch (config) # show virtual-machine host my-vm configured detail
VM 'my-vm'
   UUID:                0a177a99-f780-5951-877a-bd660e12e5db
   Text console:        enabled
   Graphics console:    enabled

   Auto-power:          last
   Boot order:          hd, cdrom
   Architecture:        x86_64
   Memory size:         512 MB
   Features:            ACPI, APIC
   Number of VCPUs:     1
      (No VCPUs pinned)

   Storage:
      IDE bus, drive 1
         Source pool:     default
         Source file:     vdisk001.img (3000 MB capacity)
         Mode:            read-write

   Interfaces:
      Interface 1
         Name:         vif1
         MAC address:  52:54:00:2F:89:69
         Model:        realtek-8139
         Bound to:     bridge 'mgmt0'
switch (config) #
```

Related Commands	
Notes	

show virtual-machine host detail

show virtual-machine host <vm-name> detail

Displays detailed status for this VM.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Role	admin	

Example

```
switch (config) # show virtual-machine host my-vm detail
VM 'my-vm'
   Status:              shut off
   UUID:                0a177a99-f780-5951-877a-bd660e12e5db
   Text console:        enabled
      Device:           N/A
   Graphics console:    enabled
      VNC display num:  N/A

   Boot order:          hd, cdrom
   Architecture:        x86_64
   Memory size:         512 MB
   Features:            ACPI, APIC
   Number of VCPUs:     1
      (State of individual VCPUs unavailable when VM is powered off)

   Storage:
      IDE bus, drive 1
         Source pool:     default
         Source file:     vdisk001.img (3000 MB capacity)
         Mode:            read-write
         Device type:     disk
         Read requests:   N/A
         Read bytes:      N/A
         Write requests:  N/A
         Write bytes:     N/A

   Interfaces:
      Interface 1
         Name:         vif1
         MAC address:  52:54:00:2F:89:69
         Model:        realtek-8139
         Bound to:     bridge 'mgmt0'
         IP address:

         RX bytes:   0                    TX bytes:   0
         RX packets: 0                    TX packets: 0
         RX errors:  0                    TX errors:  0
         RX drop:    0                    TX drop:    0
switch (config) #
```

Related Commands
Notes

show virtual-machine install

show virtual-machine host <vm-name> install

Displays status of installation of guest OS.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
	3.7.00xx	Updated Example
Role	admin	

Example	`switch (config) # show virtual-machine host my_host install`
	
	`Install status for VM 'my_host':`
	`  Install in progress, begun 9 minutes 11 seconds ago.`
	
	`  Previous install:`
	`  Completed     : 2018/09/12 14:08:45.041`
	`  Install status: FAILED`
	`  Failure reason: canceled by user`	

Related Commands		
Notes		

show virtual-machine interface

show virtual-machine host <vm-name> interface [brief | configure]

Displays full status of all interfaces for this VM.

Syntax Description	vm-name	The name of the VM.
	brief	Displays brief status of all interfaces for this VM.
	configure	Displays configuration of all interfaces for this VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
	3.7.00xx	Updated Example
Role	admin	
Example		

```
switch (config) # show virtual-machine host my-vm interface
Interface 1
   Name:         vif1
   MAC address:  52:54:00:2F:89:69
   Model:        realtek-8139
   Bound to:     bridge 'mgmt0'
   IP address:

Counters:
   RX bytes:   0                    TX bytes:   0
   RX packets: 0                    TX packets: 0
   RX errors:  0                    TX errors:  0
   RX drop:    0                    TX drop:    0
switch (config) #
```

Related Commands	
Notes	

show virtual-machine storage

show virtual-machine host <vm-name> storage

Displays statistics for attached storage.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Role	admin	

Example	`switch (config) # show virtual-machine host my-vm storage`
	`Storage for VM 'my-vm'`
	`  IDE bus, drive 1`
	`  Source pool:  default`
	`  Source file:  vdisk001.img (3000 MB capacity)`
	`  Mode:  read-write`
	`  Device type:  disk`
	`  Read requests:  N/A`
	`  Read bytes:  N/A`
	`  Write requests:  N/A`
	`  Write bytes:  N/A`
	`switch (config) #`	

Related Commands		
Notes		

4.21 Control Plane Policing

Control Plane Policing or Policies (CoPP) ensures the CPU and control plane are not over-utilized which is essential for the robustness of the switch. CoPP limits the number of control plane packets. Onyx implements several CoPP mechanisms:

- ACLs may be used to limit the rate of packets or bytes of a certain type, including L3 control packets (L2 control packets are forwarded to the CPU before the ACL)
- Policers on traffic going to the CPU – these policers are configured by Onyx and cannot be modified by the user
- IP filter tables limit the traffic to the CPU coming in from the management ports.

4.21.1 IP Table Filtering

IP table filtering is a mechanism that allows the user to apply actions to a specific control packet flow identified by a certain flow key.

This mechanism is used in order to protect switch control traffic against attacks. For example, it could allow traffic coming from a specific trusted management subnet only, block the SNMP UDP port from receiving traffic, and force ping rate to be lower than a specific threshold.

Each IP table rule is defined by key, priority, and action:

- Key – the key is a combination of physical port and layer 3 parameters (e.g. SIP, DIP, SPORT, DPORT, etc.), and other fields. Each part of the key, can be set to a specific value or masked.
- Priority – each rule in the IP table is assigned a priority, and the rule with the highest priority whose key matches the packet executes the action.
- Action – the action describes the behavior of packets which match the key. The action type may be drop, accept, rate limit, etc.

An IP table rule is bound to an IP interface that can be a management out-of-band interface, VLAN interface, or router port interface. Once bound, all traffic received (ingress rule) or transmitted (egress rule) in this direction is being verified with all bounded rules.

Once a match was found, the rule action is executed. If no match is found, the default policy of the chain shall apply.

IP table rules get a lower priority than ACL mechanism.

4.21.1.1 Configuring IP Table Filtering

Prerequisite for IPv6:

```
switch (config) # ipv6 enable
```

➢ ***To configure IPv4 table filtering:***

Step 1. Select the policy that applies to the input/output chain (default is “accept”). Run:

```
switch (config)# ip filter chain input policy drop
switch (config)# ip filter chain output policy accept
```

Step 2. Append filtering rules to the list or set a specific rule number, select a target, and (optional) any additional filter conditions. For example, run:

```
switch (config)# ip filter chain input rule append tail target rate-limit 2 protocol
udp
switch (config)# ip filter chain input rule set 2 target drop protocol icmp in-intf
mgmt1
switch (config)# ip filter chain output rule append tail target drop protocol icmp
```

Step 3. Enable IP table filtering. Run:

```
switch (config) # ip filter enable
```

Step 4. Verify IP table filtering configuration. Run:

```
switch (config) # show ip filter configured

Packet filtering for IPv4: enabled

IPv4 configuration:
  Chain 'input' Policy 'accept':
    Rule 1:
      Target      : rate-limit 2 pps
      Protocol    : udp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
      Other Filter:  -

    Rule 2:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : mgmt1 (ingress)
      State       : any
      Other Filter:  -

  Chain 'output' Policy 'accept':
    Rule 1:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
      Other Filter: -
```

4.21.1.2 Modifying IP Table Filtering

- ***To modify IP table filtering configuration:***

```
switch (config) # ip filter chain input rule modify 3 target reject-with icmp6-adm-pro-
hibited source-addr 10::0 /126
```

- ***To delete an existing IP table filtering rule:***

```
switch (config) # no ip filter chain input rule 2
```

- ***To delete all existing IP table filtering rules:***

```
switch (config) # no ip filter chain output rule all
```

- ***To insert an IP table filtering rule in a chain:***

```
switch (config) # ip filter chain input rule 2 set target drop protocol tcp dest-port
22 in-intf mgmt1
```

4.21.1.3 Rate-limit Rule Configuration

Using a rate-limit target allows to create a rule to limit the rate of certain traffic types. The limit is specified in packets per second (pps) and can be anywhere between 1-1000 pps. When enabled, the system takes the user specified rate and converts it into units of 1/10000 of a second. Therefore, any value greater than 100 can have a slight difference when the rule is displayed using the show command.

Unlike other rules which are a match type of rule, limiting packets should be followed by a rule that drops additional packets of the same “type”. Alternatively, this can be implicitly achieved by setting the chain policy to “drop” so that it drops packets not processed by matching rules. Otherwise, no effect of the rule is observed as the remaining traffic simply gets accepted.

Rate-limit is implemented with an average rate and a burst-limit. Rate values are specified in pps and take a range from 1-1000 pps. For rate values in the range 1-100, the burst value is set equal to the rate value. For rate values in the range 101-1000, the burst limit is set to 100.

4.21.2 Commands

ip filter enable
ipv6 filter enable

{ip | ipv6} filter enable
no {ip | ipv6} filter enable

Enables IP filtering.
The no form of the command disables IP filtering.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.1000
Role	admin
Example	`switch (config) # ip filter enable` `switch (config) #`
Related Commands	N/A
Notes	It is recommended to run this command only after configuring all of the IP table filter parameters.

ip filter chain policy
ipv6 filter chain policy

{ip | ipv6} filter chain <chain_name> policy {accept | drop}
no {ip | ipv6} filter chain <chain_name> policy

Configures default policy for a specific chain (if no rule matches this default policy action shall apply).
The no form of the command resets default policy for a specific chain.

Syntax Description	chain_name	Selects a chain for which to add or modify a filter: • input – input chain or ingress interfaces • output – output chain or egress interfaces
	accept	Accepts all traffic by default for this chain
	drop	Drops all traffic by default for this chain
Default	Accept for input and output chains	
Configuration Mode	config	
History	3.5.1000	
Role	admin	
Example	`switch (config) # ipv6 filter chain input policy accept` `switch (config) #`	
Related Commands	N/A	
Notes		

ip filter chain rule target
ipv6 filter chain rule target

{ip | ipv6} filter chain <chain_name> rule <oper> target <target> [<param>]
no {ip | ipv6} filter chain <chain_name> rule {<number> | all}

Inserts rule before specified rule number.
The no form of the command deletes rule for a specific chain.

Syntax Description	chain_name	A chain to which to add or modify a filter: • input – input chain or ingress interfaces • output – output chain or egress interfaces
	rule	• append tail – appends operation to the bottom of operation list • insert <oper_num> – inserts operation at specified position (existing operation at that position moves back in the list) • modify <oper_num> – modifies existing operation at specified position. Only the parameters specified in this invocation are altered; everything else is left untouched. • move <oper_num1> to <oper_num2> – moves one operation to another place in the operation list • set <oper_num> – sets operation at specified position (overwrites existing)
	target	• accept – allows the packets that match the rule into the management plane • drop – drops packets that match the rule • rate-limit – allows with rate limiting in packets per sec (PPS) • reject-with – drops the packet and replies with an ICMP error message

param	• comment <text> – specifies description string for this rule (60 chars max) • dest-addr <ip> – IP matching a specific destination address or address range. A specific IPv4 address can be provided or an entire subnet by giving an address along with netmask in dot notation or as a CIDR notation (e.g. /24). • not-dest-addr <ip> – IP not matching a specific destination address range • dest-port <port(s)> – matching a specific destination port or port range • not-dest-port <port(s)> – port not matching a specific destination port or port range • dup-delete – deletes any preexisting duplicates of this rule • in-intf – interface matching a specific inbound interface • not-in-intf <if_name> – interface not matching a specific inbound interface • out-intf <if_name> – matches a specific outbound interface • not-out-intf <if_name> – interface not matching a specific outbound interface

param4 (cont.)	• protocol <if_name> – matches a specific protocol
	  • tcp
	  • udp
	  • icmp
	  • all
	• not-protocol <protocol> – does not match a specific protocol
	  • tcp
	  • udp
	  • icmp
	  • all
	• source-addr <ip> – matches a specific source address range
	• not-source-addr <ip> – does not match a specific source address range
	• source-port <port(s)> – matches a specific source port or port range
	• not-source-port <port(s)> – does not match a specific source port or port range
	• state – matches packets in a particular state. Possible values:
	  • established – packet associated with an established connection which has seen traffic in both directions
	  • related – packet that starts a new connection but is related to an existing connection
	  • new – packet that starts a new, unrelated connection
	  • A combination can be entered separated by commas

Default	N/A
Configuration Mode	config
History	3.5.1000
Role	admin
Example	`switch (config) # ipv6 filter enable chain input rule append tail target drop state related protocol all dup-delete` `switch (config) #`
Related Commands	N/A
Notes	• The source and destination ports may each be either a single number, or a range specified as “<low>-<high>”. For example: “10-20” would specify ports 10 through 20 (inclusive). • The port parameter only works in conjunction with TCP and UDP. • Setting a “positive” rule removes any corresponding “not-” rules, and vice-versa • The “state” parameter is a classification of the packet relative to existing connections • If TCP or UDP are selected for the “protocol” parameter, source and/or destination ports may be specified. If ICMP is selected, these options are either ignored, or an error is produced.

show ip filter

	show ip filter Displays IPv4 filtering state.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ip filter`
	
	`Packet filtering for IPv4: enabled`
	
	`Active IPv4 filtering rules (omitting any not from configuration):`
	`  Chain 'input' Policy 'accept':`
	`  Rule 1:`
	`  Target  : accept`
	`  Protocol  : all`
	`  Source  : all`
	`  Destination : 1.1.1.0/24`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`  Rule 1:`
	`  Target  : reject-with icmp-net-unreachable`
	`  Protocol  : tcp`
	`  Source  : all`
	`  Destination : all`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

show ip filter all

	show ip filter all Displays IPv4 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ip filter all`
	
	`Packet filtering for IPv4: enabled`
	
	`All active IPv4 filtering rules:`
	`  Chain 'input' Policy 'accept':`
	`  Rule 1:`
	`  Target  : accept`
	`  Protocol  : all`
	`  Source  : all`
	`  Destination : 1.1.1.0/24`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`  Rule 1:`
	`  Target  : reject-with icmp-net-unreachable`
	`  Protocol  : tcp`
	`  Source  : all`
	`  Destination : all`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

show ip filter configured

	show ip filter configured Displays IPv4 filtering configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ip filter configured`
	
	`Packet filtering for IPv4: enabled`
	
	`IPv4 configuration:`
	`  Chain 'input' Policy 'accept':`
	`    Rule 1:`
	`      Target      : accept`
	`      Protocol    : all`
	`      Source      : all`
	`      Destination : 1.1.1.0/24`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`    Rule 1:`
	`      Target      : reject-with icmp-net-unreachable`
	`      Protocol    : tcp`
	`      Source      : all`
	`      Destination : all`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

show ipv6 filter

show ipv6 filter

Displays IPv6 filtering state.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ipv6 filter`
	
	`Packet filtering for IPv6: enables`
	
	`Active IPv6 filtering rules (omitting any not from configuration):`
	`  Chain 'input' Policy 'accept':`
	`    Rule 1:`
	`      Target      : accept`
	`      Protocol    : all`
	`      Source      : all`
	`      Destination : 1.1.1.0/24`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`    Rule 1:`
	`      Target      : reject-with icmp-net-unreachable`
	`      Protocol    : tcp`
	`      Source      : all`
	`      Destination : all`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

show ipv6 filter all

	show ipv6 filter all Displays IPv6 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ipv6 filter all`
	
	`Packet filtering for IPv6: enables`
	
	`All active IPv6 filtering rules:`
	`  Chain 'input' Policy 'accept':`
	`  Rule 1:`
	`  Target  : accept`
	`  Protocol  : all`
	`  Source  : all`
	`  Destination : 1.1.1.0/24`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`  Rule 1:`
	`  Target  : reject-with icmp-net-unreachable`
	`  Protocol  : tcp`
	`  Source  : all`
	`  Destination : all`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

show ipv6 filter configured

show ipv6 filter configured

Displays IPv6 filtering configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000 Updated Example
Role	admin

Example	`switch (config) # show ipv6 filter configured`
	
	`Packet filtering for IPv6: enables`
	
	`IPv6 configuration:`
	`  Chain 'input' Policy 'accept':`
	`    Rule 1:`
	`      Target      : accept`
	`      Protocol    : all`
	`      Source      : all`
	`      Destination : 1.1.1.0/24`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`    Rule 1:`
	`      Target      : reject-with icmp-net-unreachable`
	`      Protocol    : tcp`
	`      Source      : all`
	`      Destination : all`
	`      Interface   : all`
	`      State       : any`
	`      Other Filter: dest-port 1000`

Related Commands	N/A
Notes	N/A

4.22 Resource Scale

Mellanox Onyx allows dynamic allocation of internal resources so that different internal subsystems could use as much resources as are available until resource exhaustion is reached.

Internal subsystems (e.g. ACL, OF, IP router) may use internal resources according to configured allocation policy mode which, in the case of Spectrum based switch systems is Loose. Loose mode is a configuration that supports flexible user experience while providing protection to assure some protection against flooding of ARP.

Transition between modes saves configuration and reloads the system.

Table 50 presents the number of resources available for a Spectrum™ based node in loose mode.

Table 50 - Number of Resources per Node in Loose Mode

Resource	Max Resources
Number of ACL rules	5K
Number of MAC addresses	88K
Number of IPv4 neighbors	50K
Number of IPv4 UC routes	100K
Number of IPv4 MC routes	3K
Number of IPv4 (ECMP) UC routes	30K

4.22.1 Commands

show system resource table

show system resource table [<table-id>]

Displays all system resource in-use value.

Syntax Description	table-id	Displays information for a specific in-use resource table
Default	N/A	
Configuration Mode	Any command mode	
History	3.5.1000	
Role	admin	

Example	`switch (config) # show system resource table`
	`-------------------------------------`
	`Table-Id In-Use`
	`-------------------------------------`
	`acl 0`
	`ipv4-uc 1`
	`ipv4-mc 0`
	`ipv4-neigh 0`
	`ipv6-uc 0`
	`ipv6-mc 0`
	`ipv6-neigh 0`
	
	`System mode: loose`
	`Total configured entries: 1`	

Related Commands	N/A	
Notes		

4.23 Linux Dockers

Mellanox Onyx allows the user to run their own applications on a Linux docker image embedded in the switch software. The container is a pure application sandbox with resource isolation of both memory and compute from the system code/NOS.

Docker container implementation in Onyx enhances its VM support to provide a new set of capabilities:

- Network traffic access

  Docker containers are implemented in Mellanox Onyx in the same name-space as the network devices allowing the software to send and receive packets from the switch ports by opening a standard Linux socket over the network devices and using an IP address assigned to the device via the legacy management interface (e.g. JSON over HTTP).

  

  It is recommended to assign a unique port number to the Linux socket to prevent ambiguity of applications between the container and the Onyx.

- Calling the SDK interfaces

  Applications running in the docker container are able to implement a set of tools pertaining only to the container such as telemetry features within the network devices. By calling the switch SDK APIs, it can also read data that is not exposed in the Onyx user interface, or register to receive events that occur in the system (e.g. port up/down).

  

  The container implementation does not limit the container developer from calling the SDK to set parameters. However this is strongly discouraged as it may cause unexpected system behavior where the Onyx and the container application manage the same resources.

- Query the Linux tables provisioned by Onyx such as neighbor cache, routing tables, L3 interfaces attributes etc.

4.23.1 Limiting the Container's Resources

It is possible to configure multiple containers in dockers, however, they would compete for the same memory and compute resources allocated by the switch software (varies for different systems). To ensure system stability and that no random process is killed to free up memory, it is strongly recommended that all resource configurations done in the container utilize Onyx user interfaces such as JSON/SNMP and take advantage of the internal loopback interface.

4.23.1.1 Memory Resources Allocation Protocol

The Linux docker supports a hard limit to control memory resource allocation which limits the container to a given amount of user/system memory.

To set the amount of memory allocated to the container, run the following command:

```
switch (config) # docker start imagename latestver containername init memory 25 label
newlabel privileged sdk network
```

4.23.1.2 CPU Resource Allocation Protocol

Containers have unrestricted access to the host machine's CPU cycles but it is possible to set a number of constraints to limit the containers' access.

To set up limitations or regulate the containers access to CPU resources, run the following command:

```
docker start imagename latestver containername init cpus 0.2 label new_label privileged sdk
network
```

4.23.2 Commands

docker

	docker **no docker** Enables dockers then enters docker configuration context. The no form of the command disables dockers, removes configuration, and deletes all containers and docker images.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.2940
Role	admin
Example	`switch (config) # docker` `switch (config docker) #`
Related Commands	N/A
Notes	

commit

commit <container-name> <image-name> <image-version>

Creates a new image from a running container.

Syntax Description	container-name	Name of the running container to commit (limited to 180 characters)
	image-name	Name of the new image to be created
	image-version	Version of the new image to be created
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
	3.6.8008	Added new character limitation for container-name.
Role	admin	
Example	`switch (config docker) # commit mycontainer test latest`	
Related Commands	N/A	
Notes		

copy-sdk

	copy-sdk The command provides access to the switch SDK APIs giving applications running on docker access to the switch hardware.
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.4110
Role	admin
Example	`switch (config docker) # copy-sdk`
Related Commands	N/A
Notes	

remove image

remove image <image-name> <image-version>

Removes an image from the Linux docker service.

Syntax Description	image-name	Name of the new image to be deleted
	image-version	Version of the new image to be deleted
Default	N/A	
Configuration Mode	config docker	
History	3.6.3520 3.6.2940	
Role	admin	
Example	`switch (config docker) # remove image test latest`	
Related Commands	docker	
Notes		

exec

exec <container-name> <program-executable>

Executes a program within a running container.

Syntax Description	container-name	Name of the running container to commit (limited to 180 characters)
	program-executable	Linux command
Default	N/A	
Configuration Mode	config docker	
History	3.6.3520 3.6.2940	
Role	admin	
Example	`switch (config docker) # exec mycontainer "ls -la"`	
Related Commands	docker	
Notes		

label

label <label name>
no label <label name>

Creates a label which can be used as a shared storage between containers.
The no form of the command removes the label.

Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.4110
Role	admin
Example	`switch (config docker) # label new_label`
Related Commands	N/A
Notes	

load

load <image-name>

Loads an image from a TAR archive.

Syntax Description	image-name	Name of the TAR image to be loaded
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
Role	admin	
Example	`switch (config docker) # load test`	
Related Commands	docker	
Notes		

pull

pull <image-name>[:<version>]

Pulls a docker image from a docker repository.

Syntax Description	image-name	Image name Format: Name:Version If only “Name” is provided, “version” defaults to latest
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
Role	admin	

Example	`switch (config docker) # pull test`
	`Using default tag: latest`
	`latest: Pulling from library/test`
	`45a2e645736c: Pull complete`
	`Digest:`
	`sha256:c577af3197aacedf79c5a204cd7f493c8e07ffbce7f88f7600bf19c688c38799`
	`Status: Downloaded newer image for test:latest`
	`switch (config docker) #`	

Related Commands	docker	
Notes		

save

save <image-name> <image-version> <filename>

Saves an image to a TAR archive.

Syntax Description	image-name	Image name
	image-version	Image version
	filename	Name of the file in which to save the image
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
	3.6.8008	Updated command syntax
Role	admin	

Example	`switch (config docker) # save busybox latest my_image`
	
	`Saving and compressing image: busybox version: latest`
	`this could take a while...`
	
	`switch (config docker) #`	

Related Commands	docker docker load	
Notes	After the file is created, the filename gets appended a *.gz suffix.	

shutdown

	shutdown **no shutdown** Stops all docker containers, and deletes all non-auto containers. The no form of the command enables the docker Linux service and runs all configured auto-start containers
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.2940
Role	admin
Example	`switch (config docker) # no shutdown`
Related Commands	docker
Notes	

start

start <image-name> <image-version> <container-name> <starting-point> [privileged {network | sdk}] [cpus <max-cpu-resources>] [memory <max-memory>]
no start <container-name>

Starts a new container from an image.
The no form of the command stops a running docker container.

Syntax Description	image-name	Name of the new image to start
	image-version	Version of the image to start
	container-name	Name of the running container to commit (limited to 180 characters)
	privileged	• network – adds network privileges to the container (--privilege flag) • sdk – adds required mounts to use the switch SDK from the container
	starting-point	• init – persistent, start the container after boot, when system initialization is done • data-path-ready – persistent, start the container after boot, when data-path is ready to be configured • ptp-ready – persistent, start the container after boot, when protocol PTP is ready to be configured • now – start the container now, this is not persistent • now-and-data-path-ready – starts the container now and after boot, when data-path is ready to be configured • now-and-init – starts the container now and after boot, when system configuration is done
	cpus	Sets how much of the available CPU resources a container can use (e.g. “cpus 1.5” guarantees at most one and a half of the available CPUs for the container)
	memory	Sets the maximum amount of memory the container can use in MB. The minimum amount of memory to configure is 4MB.
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
	3.6.3520	Added “privileged” parameter
	3.6.8008	Added the options “now-and-data-path-ready” and “now-and-init”, new character limitation for container-name, and updated the description of the parameter “memory”

	3.7.00xx	Added “ptp-ready” option
Role	admin	

Example	`switch (config docker) # start centos latest test now`
	
	`Starting docker container. Please wait (this can take a minute)...`
	
	`switch (config) # docker start imagename latestver containername init cpus 0.2 memory 25`	

Related Commands	docker	
Notes	The no form of the command removes the container if it is not persistent.	

image upload

image upload <filename> <upload_url>

Uploads an image file to a remote host.

Syntax Description	filename	Name of file
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g. scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.6.2940	
Role	admin	
Example	`switch (config) # image upload centos.img.gz scp://username:password@192.168.10.125/var/www/html/<image_name>`	
Related Commands	N/A	
Notes		

file image upload

file image upload <filename> <upload_url>

Uploads a file to a remote host.

Syntax Description	filename	Name of file
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g. scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.6.2940	
Role	admin	
Example	`switch (config) # file image upload centos.img.gz scp://username:password@192.168.10.125/var/www/html/<image_name>`	
Related Commands	N/A	
Notes		

show docker containers

show docker containers <container_name>

Displays set parameters on containers already running, and containers planned to run in the future.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Role	admin
Example	

```
switch (config) # show docker containers

cont_example:
  image       : busybox
  version     : latest
  status      : running
  start point : data-path-ready
  cpu limit   : 0.2
  memory limit: 10m
  labels      : -
  privileges  : network, sdk

another_container:
  image       : busybox
  version     : latest
  status      : -
  start point : init
  cpu limit   : 0.2
  memory limit: 10m
  labels      : my_label
  privileges  : network, sdk

switch (config) # show docker containers cont_example

cont_example:
  image       : busybox
  version     : latest
  status      : running
  start point : data-path-ready
  cpu limit   : 0.2
  memory limit: 10m
  labels      : -
  privileges  : network, sdk
```

Related Commands	N/A
Notes	• If a container is already started, the status field displays its current status • If a container is configured to run on the next boot, the start point field displays when it will start • If there is a mismatch between the configuration of a running container and its next-boot configuration, two entries for the container are shown with both of the configurations

show docker images

	show docker images Display docker images.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940
Role	admin

Example	`switch (config) # show docker images`
	`---------------------------------------------------------------`
	`Image  Version  Created  Size`
	`---------------------------------------------------------------`
	`ubuntu  latest  Less than a secon d ago  117MB`
	`ubuntu-sdk  v1  41 seconds ago  215MB`

Related Commands	N/A
Notes	

show docker ps

	show docker ps Display docker containers.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940 Updated Example
Role	admin

Example	`switch (config) # show docker ps`
	`---------------------------------------------------------------------------`
	`Container Image:Version Created Status`
	`---------------------------------------------------------------------------`
	`my_ubuntu_app ubuntu:latest 56 seconds ago Up 50 seconds`

Related Commands	N/A
Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)

show docker labels

	show docker labels Displays docker labels.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Role	admin

Example	`switch (config) # show docker labels`
	`Storage label : label_name1`
	`  configured containers list : cont_name2`
	`  active containers list : cont_name1`
	
	`Storage label : label_name2`

Related Commands	N/A
Notes	

show docker stats

show docker stats [<name>]

Displays Linux docker statistics.

Syntax Description	name	Docker whose stats to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example

```
switch (config) # show docker stats

-------------------------------------------------------------------------------------
Container     CPU %     Memory     Memory     Memory %     Block      Block      Pids
                        Usage      Limit                   IN         OUT
-------------------------------------------------------------------------------------
test1         0.00%     104K       3.682G     0.00%        0B         0B         0
hello         0.00%     56K        3.682G     0.00%        0B         0B         0
```

Related Commands	N/A
Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)

4.24 What Just Happened (WJH)

What-just-happened is based on extended telemetry capabilities of the Spectrum family ASICs. It provides the ability to retain the last packets that where dropped from the Switch, with complete packet headers and the actual drop reason. Thus, enhancing the ability to debug network problems, identify affected flows, and decrease time-to-repair.

Obtaining What-just-happened information is done by specifically requesting the last N (up to max 1000) last dropped packets & their respective drop reasons. The information is displayed with important Ethernet, IP & L4 headers. For complete packet a PCAP file is available.

- ***To enable What Just Happened (enabled by default), run:***

```
switch (config) # what-just-happened enable
```

- ***To disable What Just Happened, run:***

```
switch (config) # no what-just-happened enable
```

- ***To automatically generate a What Just Happened PCAP file as a result of discards the following configuration is required:***

```
switch (config) # logging events interfaces enable
switch (config) # logging events interfaces interval 30
```

- ***As a result of an event the following log message will be recorded***

```
Jan  4 12:02:24 r-mgtswd-250 statsd[3138]: [statsd.NOTICE]: (StatsLog) Interface Eth1/
2: 10 packets dropped due to Rx discard packets by vlan filter
...
Jan  4 12:02:24 r-mgtswd-250 sdkd[3368]: TID 140203194615552: [sdkd.NOTICE]: WJH: What-
Just-Happened - created event pcap file: /var/opt/tms/tcpdumps/
wjh_event_2019_01_04_12_02_24.pcap
```

What Just Happened wireshark dissector enables wireshark users to analyze WJH pcap files. It displays the packets' added metadata. You may log in to the web UI and click the "Download Wireshark Plugin" button in the Status-> What Just Happened page in order to download the wireshark plugin file. After downloading the file, place it in the wireshark application, in Windows, under %APPDATA%\Wireshark\plugins. Wireshark dissector was tested on version 2.6.3.

Whenever there is a packet loss, or a critical system failure, the system will auto-generate a .pcap file under /var/opt/tms/tcpdumps. Once this is performed, WJH is enabled by default.

4.24.1 Commands

what-just-happened

what-just-happened enable
[no] what-just-happened enable

Enables/disables showing dropped packet information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
	3.7.11xx Updated example & note
Role	Admin
Example	`switch (config)# what-just-happened enable`
Related Commands	N/A
Notes	• This feature was changed from disabled to enabled by default

clear what-just-happened

clear what-just-happened

Flushes data from cache and hardware buffer.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.7.1000
Role	Admin
Example	`switch (config)# clear what-just-happened`
Related Commands	N/A
Notes	

show what-just-happened

show what-just-happened <recently | last-read> <all | forwarding> [max-packets <1-1000>] >] [no-pcap] [no-metadata] [export <file-name>]

Shows dropped packets information.

Syntax Description	recently	Shows the most recent dropped packets. Calling the "recently" option will clean the packet drop hardware buffer and cache.
	last-read	Shows the dropped packets which were previously read (using the "recently" option). This option gets the dropped packets from the cache and does not clear the hardware buffer.
	max-packets	Limit number of packets to dump.
	no-pcap	A *.pcap file with all dropped packets will be created by default. Add this flag to disable the pcap file creation and only dump the summary data to the screen.
	no-metadata	Do not add metadata to the pcap file.
	export	Change default file name.
Default	max-packets = 1000	
Configuration Mode	Any command mode	
History	3.7.11xx	Updated syntax & example
Role	Admin	

Example

```
switch (config) # show what-just-happened recently
Pcap file created: /var/opt/tms/tcpdumps/wjh_user_2019_01_01_08_02_34.pcap.
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
PktID Timestamp              sPort   dPort Size(B) VLAN    sMAC               dMAC        EthType  Src IP    Dst IP     L4 sPort L4 dPort    Drop Group    Drop Reason
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
1     2019/01/01 08:02:47.730 eth1/3   N/A         60    10    00:BB:CC:11:22:33  00:BB:CC:11:22:30      IPv4     1.2.3.4   9.9.9.9      20       80 (http)     Forwarding    VLAN filtering
```

Related Commands	N/A
Notes	

5 Ethernet Switching

5.1 Interface

Ethernet interfaces have the following physical set of configurable parameters:

- Admin state – enabling or disabling the interface
- Flow control – admin state per direction (send or receive)
- MTU (Maximum Transmission Unit) – 1500-9216 bytes
- Speed – 1/10/40/56/100GbE (depending interface type and system)
- Description – user defined string
- Module-type – the type of the module plugged in the interface

To use 100GbE QSFP interfaces as 25/10GbE (via QSA adapter), the speed must be manually set with the command “speed 25000” or “speed 10000” respectively under the interface configuration mode.

5.1.1 Break-Out Cables

The break-out cable is a unique Mellanox capability, where a single physical quad-lane QSFP port is divided into 2 dual-lane ports or 4 single-lane ports. It maximizes the flexibility of the end user to use the Mellanox switch with a combination of dual-lane, single-lane and quad-lane interfaces according to the specific requirements of its network. Certain ports cannot be split at all, and there are ports which can be split into 2 ports only (for more information please refer to your Switch System Hardware User Manual). Splitting a port changes the notation of that port from x/y to x/y/z with “x/y” indicating the previous notation of the port prior to the split and “z” indicating the number of the resulting sub-physical port (1,2 or 1,2,3,4). Each sub-physical port is then handled as an individual port. For example, splitting port 10 into 4 lanes gives the following new ports: 1/10/1, 1/10/2, 1/10/3, 1/10/4.

Figure 15: Break-Out Cable

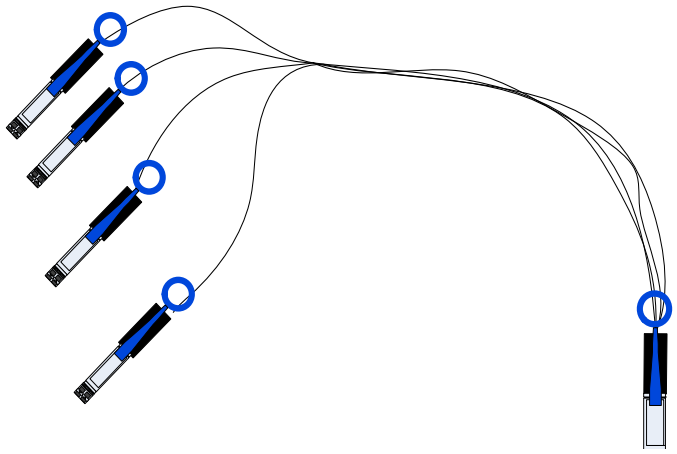

A split-4 operation results in blocking a quad-lane port in addition to the one being split. A set of hardware restrictions determine which of the ports can be split.

Specific ports can be split by using a QSFP 1X4 breakout cable to split one single-lane port into 4 lanes (4 SFP+ connectors). These 4 lanes then go, one lane to each of the 4 SFP+ connectors.

Splitting the interface deletes all configuration on that interface.

When splitting an interface's traffic into 4 data streams (four lanes) one of the other ports on the switch is disabled (unmapped).

To see the exact splitting options available per system, refer to each specific system's hardware user manual (Cabling chapter) located on the Mellanox website.

5.1.1.1 Changing the Module Type to a Split Mode

➢ ***To split an interface:***

Step 1. Shut down all the ports related to the interface. Run:

- in case of split-2, shut down the current interface only
- in case of split-4, shut down the current interface and the other interface according switch system's spec

```
switch (config) #  1/3
switch (config  1/3) # shutdown
switch (config  1/3) # exit
switch (config) #  1/4
switch (config  1/4) # shutdown
```

Step 2. Split the ports as desired. Run:

```
switch (config  1/3) # module-type qsfp-split-4
switch (config  1/3) #
```

Step 3. The following warning will be displayed:

```
The following interfaces will be unmapped: 1/3 1/4.
```

Type "Yes" when asked to confirm the split.

The <ports> field in the warning refers to the affected ports from splitting port <inf> in the applied command.

Please beware that splitting a port into 4 prevents you from accessing the splittable port, and an additional one. For example, in the procedure above, ports 3 and 4 become unaccessible.

5.1.1.2 Unsplitting a Split Port

➢ ***To unsplit a split port:***

Step 1. Shut down all of the split ports. Run:

```
switch (config  1/4/4) # shutdown
switch (config  1/4/4) # exit
switch (config) #  1/4/3
switch (config  1/4/3) # shutdown
switch (config  1/4/3) # exit
switch (config) #  1/4/2
switch (config  1/4/2) # shutdown
switch (config  1/4/2) # exit
switch (config) #  1/4/1
switch (config  1/4/1) # shutdown
```

Step 2. From the first member of the split (1/4/1), change the module-type back to QSFP. Run:

```
switch (config  1/4/1) # module-type qsfp
```

The module-type can be changed **only** from the first member of the split and **not** from the interface which has been split.

The following warning will be displayed:
`The following interfaces will be unmapped: 1/4/1 1/4/2 1/4/3 1/4/4.`

Step 3. Type "`yes`" when prompted "`Type 'yes' to confirm unsplit`."

5.1.2 56GbE Link Speed

Mellanox offers proprietary speed of 56Gb/s per Ethernet interface.

➢ ***To achieve 56GbE link speed:***

Step 1. Set the speed for the desired interface to 56GbE as follows. Run:

```
switch (config) # 1/1
switch (config 1/1) # speed 56G
switch (config 1/1) #
```

Step 2. Verify the speed is 56GbE.

```
switch (config) # show interfaces ethernet 1/1
Eth1/1
Admin state                          : Enabled
  Operational state                  : Down
  Last change in operational status: 0:00:21 ago (0 oper change)
  Boot delay time                    : 0 sec
  Description                        : N/A
  Mac address                        : 7c:fe:90:eb:52:9e
  MTU                               : 1500 bytes (Max. packet size 1522 bytes)
  Fec                                : auto
  Operational Fec                    : rs-fec
  Flow-control                       : receive off send off
  Supported speeds                   : 1G 10G 25G 40G 50G 56G 100G
  Advertised speeds                  : 56G
  Actual speed                       : Unknown
  Auto-negotiation                   : Enabled
  Width reduction mode               : Unknown
  Switchport mode                    : access
  MAC learning mode                  : Enabled
  Forwarding mode                    : inherited cut-through

  Telemetry sampling: Disabled   TCs: N/A
  Telemetry threshold: Disabled  TCs: N/A
  Telemetry threshold level: N/A

  Last clearing of "show interface" counters: 00:00:02
  60 seconds ingress rate        : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate         : 0 bits/sec, 0 bytes/sec, 0 packets/sec
```

```
Rx
    0                   packets
    0                   unicast packets
    0                   multicast packets
    0                   broadcast packets
    0                   bytes
    0                   discard packets
    0                   error packets
    0                   fcs errors
    0                   undersize packets
    0                   oversize packets
    0                   pause packets
    0                   unknown control opcode
    0                   symbol errors
    0                   discard packets by storm control

Tx
    0                   packets
    0                   unicast packets
    0                   multicast packets
    0                   broadcast packets
    0                   bytes
    0                   discard packets
    0                   error packets
    0                   hoq discard packets
```

5.1.3 Transceiver Information

Mellanox Onyx™ offers the option of viewing the transceiver information of a module or cable connected to a specific interface. The information is a set of read-only parameters burned onto the EEPROM of the transceiver by the manufacture. The parameters include identifier (connector type), cable type, speed and additional inventory attributes.

➢ ***To display transceiver information of a specific interface, run:***

```
switch (config) # show interfaces ethernet 1/1020 transceiver
Port 1/1020 state
        identifier             : QSFP+
        cable/module type      : Passive copper, unequalized
        ethernet speed and type: 56GigE
        vendor                 : Mellanox
        cable length           : 1m
        part number            : MC2207130-001
        revision               : A3
        serial number          : MT1238VS04936
switch (config) #
```

The indicated cable length is rounded up to the nearest natural number.

5.1.4 High Power Transceivers

Mellanox switch systems offer high power transceiver (LR4) support in the following ports:

Table 51 - LR4/ER4 Switch and Port Support

Transceiver			Switch OPN	Supported Ports
Speed	Protocol	Power Consumption [W]		
40GbE	LR4/ER4	3.5	SN2100/SN2410/SN2700	All ports
100GbE		3.5	SN2100/SN2410/SN2700	All ports
100GbE		4.5	SN2100	1, 2, 15, 16
			SN2410	49, 50, 55, 56
			SN2700	1, 2, 31, 32

If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when the command “show interfaces ethernet” is run.

5.1.5 Forward Error Correction

Forward Error Correction (FEC) mechanism adds extra data to the transmitted information. The receiving device uses this additional data to verify that the received data contains no errors. If the receiving side discovers errors within the received data it is able to correct some of these errors. The number or errors that can be corrected depends on the FEC algorithm and the amount of redundant data.

100GbE Mellanox-to-Mellanox Ethernet connections always enable standard Reed Solomon (RS) FEC on all cables.

If a Mellanox system is connected to a 3rd party system, then FEC is only activated if the 3rd party requests it also.

5.1.6 Commands

<slot>/<port>[/<subport>]-[<slot>/<port>[/<subport>]]

Enters the Ethernet interface or Ethernet interface range configuration mode.

Syntax Description	<slot>/<port>	Ethernet port number.
	subport	Ethernet subport number. to be used in case of split port.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.2.1100	Added range support
Role	admin	

Example	`switch (config) #  1/1`
	`switch (config  1/1) # exit`
	`switch (config) #  1/1-1/10`
	`switch (config  1/1-1/10) #`	

Related Commands	show interfaces ethernet	
Note		

boot-delay

boot-delay [<time>]
no boot-delay

Configures interface boot-delay timer.
The no form of the command returns boot-delay time to its default value.

Syntax Description	time	Boot delay time in seconds Range: 0-600
Default	0 seconds	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.2002	
Role	admin	
Example	`switch (config  1/1) # boot-delay 60`	
Related Commands	show interfaces ethernet	
Note	This command delays the interface from boot time of the interface Configuration save and system reboot is required for the configuration to take effect.	

description

description <string>
no description

Sets an interface description.
The no form of the command returns the interface description to its default value.

Syntax Description	string	40 bytes
Default	“”	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	
Example	`switch (config  1/1) # description my-interface`	
Related Commands	show interfaces ethernet	
Note		

fec-override

fec-override <fec-configuration> [force]
no fec-override <fec-configuration> [force]

Changes FEC configuration on a specific port or range of ports.
The no form of the command resets this parameter to its default value.

Syntax Description	fec-configuration	• fc-fec – FireCode FEC • no-fec – does not use FEC • rs-fec – Reed Solomon FEC
	force	Forces configuration (does not require toggling interface to take effect)
Default	Auto-FEC selection	
Configuration Mode	config interface ethernet	
History	3.5.0000	
	3.6.2002	Added force option
	3.7.1000	Example updated
Role	admin	
Example	`switch (config  1/2) # fec-override fc-fec`	
Related Commands	show interfaces ethernet	
Notes	Use this command with caution. There is no limitation in configuring non-standard FEC. It may cause the link to malfunction.	

flowcontrol

flowcontrol {receive | send} {off | on} [force]

Enables or disables IEEE 802.3x link-level flow control per direction for the specified interface.

Syntax Description	receive | send	receive - ingresses direction send - egresses direction
	off | on	on - enables IEEE 802.3x link-level flow control for the specified interface on receive or send. off - disables IEEE 802.3x link-level flow control for the specified interface on receive or send
	force	Forces command implementation.
Default	receive off, send off	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
	3.6.6102	Added Note
Role	admin	
Example	`switch (config  1/1) # flowcontrol receive off`	
Related Commands	show interfaces ethernet	
Note	To configure global pause please see Chapter 5.17.2.2, “Flowcontrol (Global pause)” on page 1107	

ip address dhcp

	ip address dhcp **no ip address dhcp** Enables DHCP on this Ethernet interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet set as router interface config interface port-channel set as router interface
History	3.4.2008
Role	admin
Example	`switch (config  1/1) # ip address dhcp`
Related Commands	show interfaces ethernet
Note	

load-interval

load-interval <time>
no load-interval

Sets the interface counter interval.
The no form of the command resets the interval to its default value.

Syntax Description	time	In seconds.
Default	300 seconds.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.3.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	
Example	`switch (config  1/1) # load-interval 30`	
Related Commands	show interfaces ethernet	
Note	This interval is used for the ingress rate and egress rate counters.	

module-type

module-type <type> [force]
no module-type <type> [force]

Splits the interface to two or four separate interfaces, or merges them back to a single interface (QSFP).
The no form of the command resets the interface to its default configuration.

Syntax Description	type	qsfp - Port runs at 40000/56000Mbps qsfp-split-2 - Port is split and runs at 2X10000Mb/s qsfp-split-4 - Port is split and runs at 4X10000Mb/s
	force	Force the split operation without asking for user confirmation.
Default	QSFP	
Configuration Mode	config interface ethernet	
History	3.1.1400	
	3.5.0000	Added note
	3.6.3640	Added note
	3.6.4006	Added note
Role	admin	

Example	`switch (config  1/4) # module-type qsfp-split-4`
	`The following interfaces will be unmapped: 1/4 1/1`
	`Type 'yes' to confirm split: yes`
	`switch (config  1/4) #`	

Related Commands		

Note	• Port cannot be split when storm-control is configured on port
	• Force command don't remove storm-control configuration. Error output:
	`% Storm control configuration must be removed from interface Eth1/2`
	• After a split port is created or deleted, the forwarding mode for each split port is set according to the global configuration
	• The affected interfaces should be disabled prior to the operation
	• In order to unsplit the interface, use the command with “qsfp”, the speed is set to 40Gb/s “module-type qsfp”
	• The following speeds are supported on the different Ethernet interface types:
	  • qsfp - 1G, 10G, 25G, 40G, 50G, 56G, 100G
	  • qsfp-split-2 - 1G, 10G, 25G, 50G
	  • qsfp-split-4 - 1G, 10G, 25G	

mtu

mtu <frame-size>

Configures the Maximum Transmission Unit (MTU) frame size for the interface.

Syntax Description	frame-size	This value may be 1500-9216 bytes
Default	1500 bytes	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	
Example	`switch (config  1/1) # mtu 9216`	
Related Commands	show interfaces ethernet	
Note		

shutdown

	shutdown **no shutdown** Disables the interface. The no form of the command enables the interface.
Syntax Description	N/A
Default	The interface is enabled.
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.1.0000 3.3.4500 Added MPO configuration mode
Role	admin
Example	`switch (config  1/1) # shutdown`
Related Commands	show interfaces ethernet
Note	

speed

speed {<value> [no-autoneg | speed_value [… speed_value]] | <auto>} [force]
no speed

Sets the speed of the interface.
The no form of the command sets the speed of the interface to its default value.

Syntax Description	value	The following speeds are available: 1G or 1000 - 1GbE 10G or 10000 - 10GbE 25G or 25000 - 25GbE 40G or 40000 - 40GbE 50G or 50000 - 50GbE 56G or 56000 - 56GbE 100G or 100000 - 100GbE auto - auto negotiates link speed (not supported on MPO or LAG interfaces)
	no-autoneg	Disallows speed auto-negotiation on the interface
	force	Forces speed change configuration
Default	Depends on the port module type, see the “Notes” section below.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.5.0000	Added 25GbE, 50GbE, and 100GbE speeds and updated notes
	3.6.6000	Added the no-autoneg parameter
Role	admin	
Example	`switch (config  1/1) # speed 40G`	

Related Commands	show interfaces ethernet
Note	• The default speed of an interface depends on its speed capabilities, an interface capable of 100GbE will have 100GbE speed by default • It is not possible to set the speed on a LAG or MPO interface • Not all interfaces support all speed options • It is not possible to set “auto” speed with the “no-autoneg” parameter • It is not possible to set “auto” speed along with specific speeds • A port with more than one speed advertised or a port configured to “auto” speed cannot be added to LAG • To change the speed of a LAG interface: 1. Remove Ethernet ports from LAG. 2. Shutdown ports. 3. Reconfigure port speed. 4. Re-enable ports. 5. Re-add ports to LAG interface.

clear counters

	clear counters Clears the interface counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ethernet config interface port-channel
History	3.1.0000 3.3.4500 Added MPO configuration mode
Role	admin
Example	`switch (config  1/1) # clear counters`
Related Commands	show interfaces ethernet
Note	This command also clears NVE counters

show interfaces counters

show interfaces <type> <id> counters [priority <prio>]

Displays the extended counters for the interface.

Syntax Description	id	Interface number: <slot>/<port>
	priority	Displays interface extended counters per priority. Range: 0-7 or “all”
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.5.1000	Added notes
	3.6.1002	Added “error packets” counter to Tx
	3.6.4006	Added extended output for storm-control
	3.6.5000	Added hoq discard packets counter
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/1 counters

Rx
  0                   packets
  0                   unicast packets
  0                   multicast packets
  0                   broadcast packets
  0                   bytes
  0                   packets of 64 bytes
  0                   packets of 65-127 bytes
  0                   packets of 128-255 bytes
  0                   packets of 256-511 bytes
  0                   packets of 512-1023 bytes
  0                   packets of 1024-1518 bytes
  0                   packets Jumbo
  0                   error packets
  0                   discard packets
  0                   hoq discard packets
  0                   fcs errors
  0                   undersize packets
  0                   oversize packets
  0                   pause packets
  0                   unknown control opcode
  0                   symbol errors
(appears only on L2 ethernet ports and port-channels supported inter-
faces)
..0                   packets
..0                   unicast packets
..0                   multicast packets
..
  0                   error packets
  0                   discard packets
  0                   discard packets by Storm Control
  0                   fcs errors
  0                   undersize packets

Tx
  0                   packets
  0                   unicast packets
  0                   multicast packets
  0                   broadcast packets
  0                   bytes
  0                   error packets
  0                   discard packets
  0                   hoq discard packets
  0                   pause packets
  0                   TX wait
  0                   TX wait useconds
  0                   queue depth TC0
  0                   queue depth TC1
  0                   queue depth TC2
  0                   queue depth TC3
  0                   queue depth TC4
  0                   queue depth TC5
  0                   queue depth TC6
  0                   queue depth TC7
```

Related Commands	
Note	Spectrum™ based systems display queue depth for TC0 - TC7

show interfaces counters discard

show interfaces <type> <id> counters discard

Displays discarded counters of the interface.

Syntax Description	id	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/9 counters discard`
	
	`Interface Eth1/9:`
	`  Rx:`
	`  0  discard packets`
	`  0  error packets`
	`  0  fcs errors`
	`  0  undersize packets`
	`  0  oversize packets`
	`  0  unknown control opcode`
	`  0  symbol errors`
	`  0  discard packets by storm control`
	`  0  general discard packets`
	`  0  policy discard packets`
	`  0  invalid tag packets`
	`  0  discard packets by vlan filter`
	
	`  Tx:`
	`  0  discard packets`
	`  0  error packets`
	`  0  hoq discard packets`
	`  0  oversize packets`
	`  0  policy discard packets`
	`  0  SLL discard packets`
	`  0  discard packets by vlan filter`
	`  0  discard packets by stp filter`
	`  0  discard packets by loopback filter`
	
	`switch (config) #`	

Related Commands		
Note		

show interfaces ethernet

show interfaces ethernet <inf>

Displays the configuration and status for the interface.

Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.1002	Added “error packets” counter to Tx, “Last change in operational status”, and “Isolation group” to output
	3.6.2002	Added “boot delay” parameters to output
	3.6.3640	Added support for “forwarding mode”
	3.6.4110	Updated Example with “Forwarding mode”
	3.6.5000	Added telemetry to output
	3.6.6000	Added output line for “auto-negotiation”
	3.6.8008	Updated Example
	3.6.8100	Updated Example
	3.7.11xx	Updated Example and notes
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/10
Eth1/10:
  Admin state                     : Enabled
  Operational state               : Up
  Last change in operational status: 0:00:47 ago (1 oper change)
  Boot delay time                 : 0 sec
  Description                     : N\A
  Mac address                     : 7c:fe:90:f5:8d:2e
  MTU                             : 1500 bytes (Maximum packet size 1522 bytes)
  Fec                             : auto
  Operational Fec                 : rs-fec
  Flow-control                    : receive off send off
  Supported speeds                : 1G 10G 25G 40G 50G 56G 100G
  Advertised speeds               : 100G
  Actual speed                    : 100G
  Auto-negotiation                : Enabled
  Width reduction mode            : Unknown
  Switchport mode                 : access
  MAC learning mode               : Enabled
  Forwarding mode                 : inherited cut-through
```

```
Telemetry sampling: Disabled   TCs: N\A
Telemetry threshold: Disabled  TCs: N\A
Telemetry threshold level: N\A

Last clearing of "show interface" counters: Never
60 seconds ingress rate                  : 232 bits/sec, 29 bytes/sec, 1 packets/sec
60 seconds egress rate                   : 8 bits/sec, 1 bytes/sec, 1 packets/sec

Rx:
  25                    packets
  0                     unicast packets
  25                    multicast packets
  0                     broadcast packets
  1600                  bytes
  0                     discard packets
  0                     error packets
  0                     fcs errors
  0                     undersize packets
  0                     oversize packets
  0                     pause packets
  0                     unknown control opcode
  0                     symbol errors
  0                     discard packets by storm control

Tx:
  3                     packets
  0                     unicast packets
  3                     multicast packets
  0                     broadcast packets
  192                   bytes
  0                     discard packets
  0                     error packets
  0                     hoq discard packets
```

Related Commands

Note

- If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when running the command “show interfaces ethernet” is run. For more information, please refer to Section 5.1.4, “High Power Transceivers,” on page 711.
- "Operational Fec" will be N/A while port is DOWN, no-fec/fc-fec/rs-fec while port is UP.

show interfaces ethernet description

show interfaces ethernet [<inf>] description

Displays the admin status and protocol status for the specified interface.

Syntax Description	inf	Interface number: <slot>/<port>.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1100	Updated Example
Role	admin	

Example	`switch (config) # show interfaces ethernet description`
	
	`Interface Admin state Operational state`
	`--------- ----------- -----------------`
	`Eth1/58 Enabled Down`
	`Eth1/59 Enabled Up`
	`Eth1/60 Enabled Down (Suspend)`
	`switch (config) # show interfaces ethernet 1/60 description`
	
	`Eth1/60`
	
	`Admin state: Enabled`
	`Operational state: Down (Suspend)`
	`switch (config) #`	

Related Commands		
Note		

show interfaces ethernet rates

show interfaces ethernet rates [<transfer-rate-unit>]

Displays the current transfer rate of the interface.

Syntax Description	transfer-rate-unit	• bytes - displays interface transfer rates in B/s dynamically (while converting to K/M/G if needed). • KB – displays interface transfer rate in Kb/s • MB – displays interface transfer rate in Mb/s • GB – displays interface transfer rate in Gb/s • bits – displays interface transfer rates in b/s dynamically (while converting to K/M/G if needed). • Kb – displays interface transfer rate in Kb/s • Mb – displays interface transfer rate in Mb/s • Gb – displays interface transfer rate in Gb/s • If no parameter is entered, transfer rate is displayed in bits.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2002 3.7.00xx	
Role	admin	
Example	`switch (config) # show interfaces ethernet rates KB`	

```
Port                        egress                          ingress
                  avg rate (KB/s)   pkts/sec        avg rate (KB/s)   pkts/sec
---------        ----------------   --------        ---------------   --------
Eth1/1                        0           0                 0.032           1
Eth1/2                        0           0                 0.032           1
Eth1/3                        0           0                     0           0
...
switch (config) #
```

Related Commands	
Note	• Added new rates to “transfer-rate-unit”

show interfaces ethernet status

show interfaces ethernet [<inf>] status

Displays the status, speed and negotiation mode of the specified interface.

Syntax Description	inf	Interface number: <slot>/<port>.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1100	Updated Example
Role	admin	
Example	`switch (config) # show interfaces ethernet status`	

```
Port                Operational state          Speed                Negotiation
----                -----------------          -----                -----------
Eth1/58             Down                       40 Gbps           No-Negotiation
Eth1/59             Up                         40 Gbps           No-Negotiation
Eth1/60             Down (Suspend)             40 Gbps           No-Negotiation
switch (config) #
```

Related Commands	
Note	

show interfaces ethernet transceiver

show interfaces ethernet [<inf>] transceiver

Displays the transceiver info.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 transceiver`
	`Port 1/1 state`
	`identifier : QSFP+`
	`cable/module type : Optical cable/module`
	`ethernet speed and type: 40GBASE - SR4`
	`vendor : Mellanox`
	`cable_length : 50 m`
	`part number : MC2210411-SR4`
	`revision : A1`
	`serial number : TT1151-00006`
	`switch (config) #`	

Related Commands		
Note	• For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage in Mellanox.com: http://www.mellanox.com/page/cables?mtag=cable_overview. • If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”.	

show interfaces ethernet transceiver brief

show interfaces ethernet [<inf>] transceiver brief

Display brief transceiver info for this interface.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 transceiver brief`
	`show interfaces ethernet transceiver brief`
	
	`---`
	`Interface Identifier Vendor PN SN Rev`
	`---`
	`Eth1/1`
	`Eth1/2 QSFP+ Mellanox MCP1600-E00A MT1710VS06916 A3`
	`Eth1/3 QSFP+ Mellanox MCP1600-E00A MT1710VS06929 A3`
	`Eth1/4 QSFP+ Mellanox MCP1600-E00A MT1710VS06953 A3`
	`Eth1/5 QSFP+ Mellanox MCP1600-E00A MT1710VS06923 A3`	

Related Commands		
Note	• For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage in Mellanox.com: http://www.mellanox.com/page/cables?mtag=cable_overview. • If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”.	

show interfaces ethernet transceiver counters

show interfaces ethernet [<inf>] transceiver counters

Displays PHY counters.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 transceiver counters`
	
	`Rx`
	`phy received bits 17725862707200`
	`phy symbol errors 0`
	`phy corrected bits 0`	

Related Commands		
Note	• The counter “phy received bits” provides information on the total amount of traffic received and can be used to estimate the ratio of error traffic • The counter “phy symbol errors” provides information on the error traffic that was not corrected because the FEC algorithm could not do it or because FEC was not active on this interface • The counter “phy corrected bits” provides the number of corrected bits by the active FEC mode (RS/FC)	

show interfaces ethernet transceiver counters details

show interfaces ethernet [<inf>] transceiver counters

Displays all PHY counters.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Role	admin	
Example	`switch (config) # show interfaces ethernet 1/1 transceiver counters details`	
Related Commands		
Note	The number of lanes displayed depends on interface splitter ratio (4-way-split – each split has only 1 lane; 2-way-split – each split has 2 lanes)	

```
-------------------------------------------------
Phy counters
-------------------------------------------------
Symbol errors                   0
Sync headers errors             0
Edpl/bip errors lane0           0
Edpl/bip errors lane1           0
Edpl/bip errors lane2           0
Edpl/bip errors lane3           0
FC corrected blocks lane0       0
FC corrected blocks lane1       0
FC corrected blocks lane2       0
FC corrected blocks lane3       0
FC uncorrectable blocks lane0   0
FC uncorrectable blocks lane1   0
FC uncorrectable blocks lane2   0
FC uncorrectable blocks lane3   0
RS corrected blocks             0
RS uncorrectable blocks         0
RS no errors blocks             1130552748
RS single error blocks          0
RS corrected symbols total      0
RS corrected symbols lane0      0
RS corrected symbols lane1      0
RS corrected symbols lane2      0
RS corrected symbols lane3      0
Link down events                0
Successful recovery events      0
Time since last clear           176127
```

show interfaces ethernet transceiver diagnostics

show interfaces ethernet [<inf>] transceiver diagnostics

Displays cable channel monitoring and diagnostics info for this interface. Tx and Rx power are reported in mW and dBm units.

Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2002	
	3.6.4006	Updated Example to report Tx and Rx power in mW and dBm units.
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/5 transceiver diagnostics

Port 1/5 transceiver diagnostic data:
  Temperature (-127C to +127C):
    Temperature          : 26 C
    Hi Temp Alarm Thresh : 80 C
    Low Temp Alarm Thresh: -10 C
    Temperature Alarm    : None

  Voltage ( 0 to 6.5535 V):
    Voltage              : 3.28980 V
    Hi Volt Alarm Thresh : 3.50000 V
    Low Volt Alarm Thresh: 3.10000 V
    Voltage Alarm        : None

  Tx Bias Current ( 0 to 131 mA):
    Ch1 Tx Current          : 6.60000 mA
    Ch2 Tx Current          : 6.60000 mA
    Ch3 Tx Current          : 6.60000 mA
    Ch4 Tx Current          : 6.60000 mA
    Hi Tx Crnt Alarm Thresh : 8.50000 mA
    Low Tx Crnt Alarm Thresh: 5.49200 mA
    Ch1 Tx Current Alarm    : None
    Ch2 Tx Current Alarm    : None
    Ch3 Tx Current Alarm    : None
    Ch4 Tx Current Alarm    : None

  Tx Power ( 0 mW to 6.5535 mW / 8.1647 dBm):
    Ch1 Tx Power             : 1.01420 mW / 0.06124 dBm
    Ch2 Tx Power             : 0.96740 mW / -0.14394 dBm
    Ch3 Tx Power             : 0.96730 mW / -0.14439 dBm
    Ch4 Tx Power             : 0.96050 mW / -0.17503 dBm
    Hi Tx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm
    Low Tx Power Alarm Thresh: 0.07240 mW / -11.40261 dBm
    Ch1 Tx Power Alarm       : None
    Ch2 Tx Power Alarm       : None
    Ch3 Tx Power Alarm       : None
    Ch4 Tx Power Alarm       : None

  Rx Power ( 0 mW to 6.5535 mW / 8.1647 dBm):
    Ch1 Rx Power             : 0.99160 mW / -0.03663 dBm
    Ch2 Rx Power             : 1.06080 mW / 0.25633 dBm
    Ch3 Rx Power             : 1.09810 mW / 0.40642 dBm
    Ch4 Rx Power             : 0.97500 mW / -0.10995 dBm
    Hi Rx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm
    Low Rx Power Alarm Thresh: 0.04670 mW / -13.30683 dBm
    Ch1 Rx Power Alarm       : None
    Ch2 Rx Power Alarm       : None
    Ch3 Rx Power Alarm       : None
    Ch4 Rx Power Alarm       : None

  Vendor Date Code (dd-mm-yyyy): 07-11-2016
```

Related Commands

Note This example is for a QSFP transceiver

show interfaces ethernet transceiver raw

show interfaces ethernet [<inf>] transceiver raw

Displays cable info for this interface.

Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Role	admin	
Example		

```
switch (config) # show interfaces ethernet 1/7 transceiver raw
Port 1/7 raw transceiver data:

I2C Address 0x50, Page 0, 0:255:
  0000  0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0030  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0040  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0050  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0060  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0070  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0080  0d 00 23 08 00 00 00 00 00 00 00 05 8d 00 00 00  ..#.............
  0090  00 00 01 a0 4d 65 6c 6c 61 6e 6f 78 20 20 20 20  ....Mellanox
  00a0  20 20 20 20 0f 00 02 c9 4d 43 32 32 30 37 31 33      ....MC220713
  00b0  30 2d 30 30 41 20 20 20 41 33 02 03 05 00 46 66  0-00A   A3....Ff
  00c0  00 00 00 00 4d 54 31 32 32 37 56 53 30 30 36 34  ....MT1227VS0064
  00d0  32 20 20 20 31 32 30 37 30 38 20 20 00 00 00 e4  2   120708  ....
  00e0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00f0  00 00 00 00 00 00 00 00 00 00 02 00 00 30 00 00
I2C Address 0x50, Pages 1, 128:255:
  0080  0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0090  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00a0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00b0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00c0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00d0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00e0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00f0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
...
```

Related Commands	
Note	

show interfaces status

	show interfaces status Displays the configuration and status for the interface
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4006
Role	admin

Example	`switch (config) # show interfaces status`
	`----------------------------------------------------------------------`
	`Port  Oper State  Admin  Speed  Description`
	`----------------------------------------------------------------------`
	`mgmt0  Down  Enabled  1000Mb/s (auto)  -`
	`mgmt1  Down  Enabled  UNKNOWN  -`
	`Eth1/1  Down  Enabled  100 Gbps  -`
	`Eth1/2  Up  Enabled  100 Gbps  -`
	`Eth1/3  Up  Enabled  100 Gbps  -`
	`Eth1/4  Up  Enabled  100 Gbps  -`
	`Eth1/5  Up  Enabled  100 Gbps  -`
	`Eth1/6  Down  Enabled  100 Gbps  -`
	`Eth1/7  Down  Enabled  100 Gbps  -`
	`Eth1/8  Down  Enabled  100 Gbps  -`
	`Eth1/9  Down  Enabled  100 Gbps  -`
	`Eth1/10  Up  Enabled  100 Gbps  -`
	`Eth1/11  Up  Enabled  100 Gbps  -`
	`Eth1/12  Up  Enabled  100 Gbps  -`
	`Eth1/13  Up  Enabled  100 Gbps  -`
	`Eth1/14  Down  Enabled  100 Gbps  -`
	`Eth1/15  Up  Enabled  100 Gbps  -`
	`Eth1/16  Up  Enabled  100 Gbps  -`
	`Eth1/17  Down  Enabled  100 Gbps  -`
	`Eth1/18  Down  Enabled  100 Gbps  -`
	`...`

Related Commands	
Note	If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when running the command “show interfaces ethernet” is run. For more information, please refer to Section 5.1.4, “High Power Transceivers,” on page 711.

5.2 Interface Isolation

Interface isolation provides the ability to group interfaces in sets where traffic from each port is isolated from other interfaces in the group. The isolated interfaces in the group, however, are able to communicate with the interface marked as privileged.

5.2.1 Configuring Isolated Interfaces

Figure 16: Interface Isolation Example

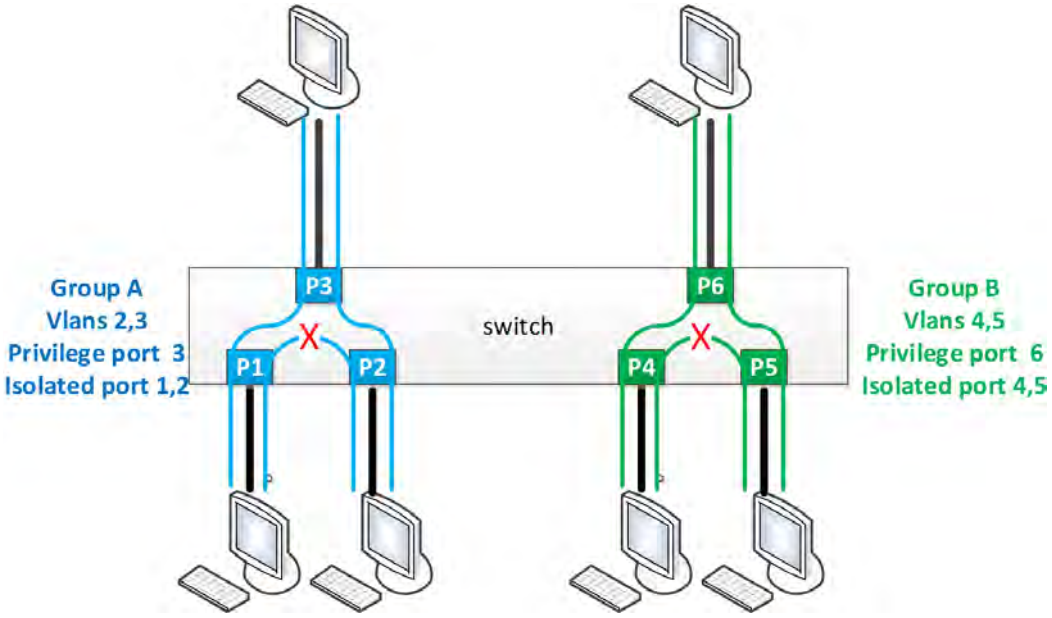

➢ ***To configure isolated interfaces:***

Step 1. Create the VLANs to be used. Run:

```
switch (config) # vlan 2-5
(config vlan 2-5) # exit
```

Step 2. Unlock isolation interface protocol. Run:

```
switch (config) # protocol isolation-group
```

Step 3. Create isolation Group A. Run:

```
switch (config) # isolation-group GroupA
```

Step 4. Assign VLANs 2 and 3 to isolation Group A. Run:

```
(config isolation-group GroupA) # vlan 2-3
(config isolation-group GroupA) # exit
```

Step 5. Create isolation Group B. Run:

```
switch (config) # isolation-group GroupB
```

Step 6. Assign VLANs 4 and 5 to isolation Group B. Run:

```
(config isolation-group GroupB) # vlan 4-5
(config isolation-group GroupB) # exit
```

Step 7. Set Ethernet interfaces 1-3 to access for VLAN 3. Run:

```
(config) # interface ethernet 1/1 switchport access vlan 3
(config) # interface ethernet 1/2 switchport access vlan 3
(config) # interface ethernet 1/3 switchport access vlan 3
```

Step 8. Isolate Ethernet interfaces 1 and 2 and set Ethernet interfaces 3 as privileged. Run:

```
(config) # interface ethernet 1/1-1/2 isolation-group GroupA mode isolated
(config) # interface ethernet 1/3 isolation-group GroupA mode privileged
```

Step 9. Enable isolation Group A. Run:

```
(config) # isolation-group GroupA no shutdown
```

Step 10. Set Ethernet interfaces 4-6 to trunk. Run:

```
(config) # interface ethernet 1/4 switchport mode trunk
(config) # interface ethernet 1/5 switchport mode trunk
(config) # interface ethernet 1/6 switchport mode trunk
```

Step 11. Isolate Ethernet interfaces 4 and 5 and set Ethernet interfaces 6 as privileged. Run:

```
(config) # interface ethernet 1/4-1/5 isolation-group GroupA mode isolated
(config) # interface ethernet 1/6 isolation-group GroupA mode privileged
```

Step 12. Enable isolation Group B. Run:

```
(config) # isolation-group GroupB no shutdown
```

Step 13. Verify configuration. Run:

```
(config) # show isolation-group
Isolation group: GroupA
State:           Enabled
VLANs:           2, 3
Privileged port: Eth1/3
Isolated ports:  Eth1/1, Eth1/2

Isolation group: GroupB
State:           Enabled
VLANs:           4, 5
Privileged port: Eth1/6
Isolated ports:  Eth1/4, Eth1/5
```

5.2.2 Commands

protocol isolation-group

	protocol isolation-group **no protocol isolation-group** Enables interface isolation and unlocks further isolation-group commands. The no form of the command disables interface isolation and locks other isolation-group commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.1002
Role	admin
Example	`switch (config) # protocol isolation-group`
Related Commands	
Note	• MLAG must be disabled before enabling interface isolation • When disabled, all configuration is lost

isolation-group

	isolation-group <name> **no isolation-group <name>** Creates isolation group. The no form of the command deletes isolation group.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.1002
Role	admin
Example	`switch (config) # isolation-group mygroup`
Related Commands	protocol isolation-group
Note	• The no form of this command deletes the isolation group, removes its attached ports, and the VLANs from the group • Up to 64 isolation groups can be created

shutdown

	shutdown **no shutdown** Enables isolation group. The no form of the command disables isolation group.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config isolation group
History	3.6.1002
Role	admin
Example	`switch (config isolation-group mygroup) # no shutdown`
Related Commands	protocol isolation-group isolation-group
Note	Enabling isolation groups fails if there are VLANs with ports both inside and outside the group.

vlan

	vlan <vid> **no vlan <vid>** Adds a VLAN to isolation group. The no form of the command removes a VLAN from an isolation group.
Syntax Description	N/A
Default	N/A
Configuration Mode	config isolation group
History	3.6.1002
Role	admin
Example	`switch (config isolation-group mygroup) # vlan 10`
Related Commands	protocol isolation-group isolation-group
Note	• Enabling isolation groups fails if there are VLANs with ports both inside and outside the group • The VLAN must be created before running this command • All interfaces in the VLAN must be attached to only this isolation group • The VLAN added cannot have a respective VLAN interface

isolation-group mode

isolation-group <name> mode {isolated | privileged}
no isolation-group <name> mode {isolated | privileged}

Adds a VLAN to isolation group.
The no form of the command removes a VLAN from an isolation group.

Syntax Description	name	The isolation group name
	isolated	Configures this interface as isolated
	privileged	Configures this interface as privileged
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/2) # isolation-group mygroup mode privileged`	
Related Commands		
Note	• Enabling isolation groups fails if there are VLANs with ports both inside and outside the group • The VLAN must be created before running this command • All interfaces in the VLAN must be attached to only this isolation group • The VLAN added cannot have a respective VLAN interface	

show isolation-group

show isolation-group <name>

Displays isolation group information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.1002 3.6.5000 Updated Example
Role	admin

Example	`switch (config) # show isolation-group mygroup`
	`Isolation group 1:`
	`  State:  Disabled`
	`  VLANs:  N/A`
	`  Privileged port:  N/A`
	`  Isolated ports:  N/A`

Related Commands	
Note	

5.3 Link Aggregation Group (LAG)

Link Aggregation protocol describes a network operation in which several same speed links are combined into a single logical entity with the accumulated bandwidth of the originating ports. LAG groups exchange Lag Aggregation Control Protocol (LACP) packets in order to align the functionality between both endpoints of the LAG. To equally send traffic on all LAG links, the switch uses a hash function which can use a set of attributes as key to the hash function.

As many as 16 physical ports can be aggregated on a single LAG.

5.3.1 Configuring Static Link Aggregation Group (LAG)

➢ ***To configure a static LAG:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a port-channel entity. Run:

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

Step 4. Change back to config mode.

```
switch (config interface port-channel 1) # exit
switch (config) #
```

Step 5. Add a physical port to the port-channel. Run:

```
switch (config  1/4) # channel-group 1 mode on
switch (config  1/4) #
```

If the physical port is operationally up, this port becomes an active member of the aggregation. Consequently, it becomes able to convey traffic.

5.3.2 Configuring Link Aggregation Control Protocol (LACP)

➢ ***To configure LACP:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a port-channel entity. Run:

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

Step 4. Change back to config mode. Run:

```
switch (config interface port-channel 1) # exit
switch (config) #
```

Step 5. Enable LACP in the switch. Run:

```
switch (config) # lacp
switch (config) #
```

Step 6. Add a physical port to the port-channel. Run:

```
switch (config  1/4) # channel-group 1 mode active/passive
switch (config  1/4) #
```

5.3.3 Commands

interface port-channel

interface port-channel <1-4096>[-<2-4096>]
no interface port-channel <1-4096>[-<2-4096>]

Creates a LAG and enters the LAG configuration mode. There is an option to create a range of LAG interfaces.
The no form of the command deletes the LAG, or range of LAGs.

Syntax Description	1-4096 / 2-4096	LAG number
Default	N/A	
Configuration Mode	config	
History	3.1.1400	
	3.2.1100	Added range support
	3.4.0000	Added note
	3.6.3640	Added note
Role	admin	

Example	`switch (config)# interface port-channel 1`
	`switch (config interface port-channel 1) # exit`
	`switch (config)# interface port-channel 1-10`
	`switch (config interface port-channel 1-10) #`	

Related Commands		
Note	• If a LAG is also an IPL, attempting to delete it without first deleting the IPL is rejected by the management. • LAG have forwarding mode in accordance with the global configuration	

lacp

	lacp **no lacp** Enables LACP in the switch. The no form of the command disables LACP in the switch.
Syntax Description	N/A
Default	LACP is disabled.
Configuration Mode	config
History	3.1.1400
Role	admin
Example	`switch (config)# lacp` `switch (config)#`
Related Commands	
Note	

lacp system-priority

lacp system-priority <1-65535>
no lacp system-priority

Configures the LACP system priority.
The no form of the command sets the LACP system-priority to default.

Syntax Description	1-65535	LACP system-priority.
Default	32768	
Configuration Mode	config	
History	3.1.1400	
Role	admin	

Example	`switch (config)# lacp system-priority 1`
	`switch (config)# show lacp interfaces port-channel`
	`Port-channel Module Admin Status is enabled`
	`Port-channel System Identifier is 00:02:c9:5c:61:70`
	`LACP System Priority: 3`
	`switch (config)#`	

Related Commands		
Note		

lacp (interface)

lacp {rate fast | port-priority <1-65535>}
no lacp {rate fast | port-priority}

Configures the LACP interface parameters.
The no form of the command sets the LACP interface configuration to default.

Syntax Description	rate fast	Sets LACP PDUs on the port to be in fast (1 second) or slow rate. (30 seconds).
	1-65535	LACP port-priority.
Default	rate - slow (30 seconds) port-priority 32768	
Configuration Mode	config	
History	3.1.1400	
Role	admin	
Example	`switch (config  1/7)# lacp rate fast`	
Related Commands		
Note	Configuring LACP rate (fast or slow) will configure the peer port to send (fast or slow), it does not make any affect on the local port LACP rate.	

port-channel load-balance ethernet

port-channel load-balance ethernet <method>
no port-channel load-balance ethernet <method>

Configures the port-channel load balancing distribution function method.
The no form of the command sets the distribution function method to default.

Syntax Description	method	Possible load balance methods: • destination-ip • destination-mac • destination-port • source-destination-ip • source-destination-mac • source-destination-port • source-ip • source-mac • source-port
Default	source-destination-mac	
Configuration Mode	config	
History	3.1.1400	
Role	admin	
Example	`switch (config)# port-channel load-balance ethernet destination-ip source-port source-mac`	
Related Commands		
Note	Several load balance methods can be configured (refer to the example)	

channel-group

channel-group <1-4096> [mode {on | active | passive}]
no channel-group

Assigns and configures a physical interface to a port channel.
The no form of the command removes a physical interface from the port-channel.

Syntax Description	1-4096	The port channel number.
	mode on	Static assignment the port to LAG. LACP will not be enabled on this port.
	mode active/passive	Dynamic assignment of the port to LAG. LACP will be enabled in either passive or active mode.
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.1.1400	
	3.4.0008	Added a note
	3.6.3640	Added a note
	3.6.4006	Added a note
Role	admin	
Example	`switch (config  1/7)# channel-group 1 mode active`	
Related Commands	show interfaces port-channel summary show interfaces port-channel compatibility-parameters show lacp interfaces ethernet	

Note	• Setting the mode to active/passive is possible only in LACP is enabled.
	• The first port in the LAG decide if the LAG will be static (“on”) or LACP (“active” , “pasive”).
	• All the ports in the LAG must have the same configuration, determines by the first port added to the LAG. The port with a different configuration will be rejected, for the list of dependencies refer to ‘show interfaces port-channel compatibility-parameters’
	• A physical port may only be part of one channel-group
	• Added support to check if the forwarding mode of the interface is the same as the forwarding mode of LAG. Error output:
	`% Channel-group and Ethernet port have different port forwarding mode configuration`
	• Port cannot be added to port-channel when storm-control is configured on port. Error output:
	`% Interface * has storm control configuration and can't be added to LAG`	

lacp-individual enable

lacp-individual enable [force]
no lacp-individual enable [force]

Configures the LAG to act with LACP-individual capabilities.
The no form of the command disables the LACP-individual capability.

Syntax Description	force	Toggles the interface after enabling LACP-individual.
Default	N/A	
Configuration Mode	config interface port-channel	
History	3.4.1100	
Role	admin	
Example	`switch (config interface port-channel 10)# lacp-individual enable force`	
Related Commands		
Note	If a switch is connected via LAG to a host without LACP capability, running this command on that LAG allows a member port (with the lowest numerical priority value), acting as an individual, to communicate with the host.	

ip address dhcp

	ip address dhcp **no ip address dhcp** Enables DHCP on this LAG interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface port-channel set as router interface
History	3.4.2008
Role	admin
Example	`switch (config interface port channel 10) # ip address dhcp` `switch (config interface port channel 10) #`
Related Commands	interface port-channel show interface port-channel
Note	

show lacp counters

	show lacp counters Displays the LACP PDUs counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.6.6000 Updated Example
Role	admin

Example	`switch (config)# show lacp counters`
	`VRF Name: default`
	
	`Port-channel 5:`
	`---------------------------------------------------------------------------------`
	`LACPDUs Marker Marker Marker Rsp Marker Rsp LACPDUs LACPDUs Illegal Unknown`
	`Port Sent Recv Sent Recv Sent Recv`
	`---------------------------------------------------------------------------------`
	`1/12 0 0 0 0 0 0 0 0`
	`1/11 0 0 0 0 0 0 0 0`
	`1/10 0 0 0 0 0 0 0 0`

Related Commands	
Note	

show lacp interfaces ethernet

show lacp <inf>

Displays the LACP interface configuration and status.

Syntax Description	inf Interface number, for example “1/1”.
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.6.6102
Role	admin

Example	`switch (config) # show lacp interfaces ethernet 1/1`
	`Port: 1/1`
	`Port State: Down`
	`Channel Group: 1`
	`Pseudo port-channel: Po1`
	`LACP port-priority: 32768`
	`LACP Rate: Slow`
	`LACP Activity: Active`
	`LACP Timeout: Short`
	`Aggregation State: Aggregation, Defaulted,`
	
	`---------------------------------------------------------------`
	`LACP Port Admin Oper Port Port`
	`Port State Priority Key Key Number State`
	`---------------------------------------------------------------`
	`1/1 Down 32768 13826 13826 0x1 0x0`

Related Commands	
Note	

show lacp interfaces neighbor

	show lacp interfaces neighbor Displays the LACP interface neighbor status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.4.0000 Updated Example
Role	admin

Example

```
switch (config) # show lacp interfaces neighbor
Flags:
A - Device is in Active mode
P - Device is in Passive mode

Channel group 1 neighbors

Port 1/4
----------
Partner System ID               : 00:00:00:00:00:00
Flags                           : A
LACP Partner Port Priority      : 0
LACP Partner Oper Key           : 0
LACP Partner Port State         : 0x0

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing

MLAG channel group 25 neighbors

Port 1/49
----------
Partner System ID               : 00:02:c9:fa:c4:c0
Flags                           : A
LACP Partner Port Priority      : 255
LACP Partner Oper Key           : 33
LACP Partner Port State         : 0xbc

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing,

MLAG channel group 28 neighbors

Port 1/51
----------
Partner System ID               : f4:52:14:10:d8:f1
Flags                           : A
LACP Partner Port Priority      : 255
LACP Partner Oper Key           : 33
LACP Partner Port State         : 0xbc

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing,

switch (config) #
```

Related Commands

Note

show lacp

	show lacp Displays the LACP global parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.0000
Role	admin

Example	`switch (config) # show lacp`
	`Port-channel Module Admin Status is enabled`
	`switch (config) #`

Related Commands	
Note	

show lacp interfaces system-identifier

	show lacp interfaces {mlag-port-channel | port-channel} <instance> system-identifier Displays the system identifier of LACP.
Syntax Description	instance LAG or MLAG instance.
Default	N/A
Configuration Mode	Any command mode
History	3.4.0000
Role	admin

Example	`switch (config)# show lacp interfaces port-channel 2 system-identifier`
	`Priority: 12345`
	`MAC: 00:02:C9:AC:2A:60`
	`switch (config)#`

Related Commands	
Note	

show interfaces port-channel

show interfaces port-channel <port-channel>

Displays LAG configuration properties.

Syntax Description	port-channel	LAG interface whose properties to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4000	
	3.4.1100	Updated Example
	3.6.1002	Added “error packets” counter to Tx
	3.6.5000	Updated Example with telemetry
	3.6.8008	Updated Example
	3.7.1000	Updated Example
Role	admin	

Example

```
switch (config) # show interfaces port-channel 1

Po1:
  Admin state         : Enabled
  Operational state   : Down
  Description         : N/A
  Mac address         : 24:8A:07:83:30:C8
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : N/A
  Width reduction mode: Not supported
  DHCP client         : Disabled
  Autoconfig          : Disabled

  IPv4 address:
    192.168.100.254/24 [primary]
    192.168.110.254/24

  Broadcast address:
    192.168.100.255 [primary]
    192.168.110.255

  IPv6 address:
    6000::1/64 [primary]
    7000::1/64

  Arp responder  : Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                    : 0 bits/sec, 0 bytes/sec, 0 packets/sec
```

```
Rx:
   0                    packets
   0                    unicast packets
   0                    multicast packets
   0                    broadcast packets
   0                    bytes
   0                    discard packets
   0                    error packets
   0                    fcs errors
   0                    undersize packets
   0                    oversize packets
   0                    pause packets
   0                    unknown control opcode
   0                    symbol errors

 Tx:
   0                    packets
   0                    unicast packets
   0                    multicast packets
   0                    broadcast packets
   0                    bytes
   0                    discard packets
   0                    error packets
   0                    hoq discard packets
```

Related Commands

Note

show interfaces port-channel counters

show interfaces port-channel <port-channel> counters

Displays the extended counters for the interface.

Syntax Description	port-channel	LAG interface whose properties to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Role	admin	

Example	`switch (config) # show interfaces port-channel 3 counters`
	
	`Rx`
	`  0  packets`
	`  0  unicast packets`
	`  0  multicast packets`
	`  0  broadcast packets`
	`  0  bytes`
	`  0  packets of 64 bytes`
	`  0  packets of 65-127 bytes`
	`  0  packets of 128-255 bytes`
	`  0  packets of 256-511 bytes`
	`  0  packets of 512-1023 bytes`
	`  0  packets of 1024-1518 bytes`
	`  0  packets Jumbo`
	`  0  error packets`
	`  0  discard packets`
	`  0  fcs errors`
	`  0  undersize packets`
	`  0  oversize packets`
	`  0  pause packets`
	`  0  unknown control opcode`
	`  0  symbol errors`
	
	`Tx`
	`  1000000  packets`
	`  0  unicast packets`
	`  1000000  multicast packets`
	`  0  broadcast packets`
	`  1505000000  bytes`
	`  1000000  error packets`
	`  0  discard packets`
	`  0  pause packets`
	`switch (config) #`	

Related Commands		
Note		

show interfaces port-channel compatibility-parameters

	show interfaces port-channel compatibility-parameters Displays LAG parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4000 3.6.3640 — Added “forwarding mode” as compatibility parameter to output 3.6.6000 — Updated Example 3.6.8008 — Updated Example
Role	admin

Example	`switch (config) # show interfaces port-channel compatibility-parameters`
	
	`Compatibility-parameters:`
	`  * Port-mode`
	`  * Speed`
	`  * MTU`
	`  * Forwarding mode`
	`  * Flow Control`
	`  * Access VLAN`
	`  * Allowed VLAN list`
	`  * Flowcontrol & PFC`
	`  * Channel-group mode`
	`  * QoS parameters`
	`  * MAC learning disable`
	
	`Static configuration on the port should be removed:`
	`  * ACL port binding`
	`  * Static mrouter`
	`  * sflow`
	`  * OpenFlow`
	`  * port mirroring local analyzer port`
	`  * Static mac address`

Related Commands	
Note	

show interfaces port-channel load-balance

show interfaces port-channel load-balance

Displays the type of load-balancing in use for LAGs.

Syntax Description	N/A	N/A
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4000	
Role	admin	
Example	`switch (config) # show interfaces port-channel load-balance` `source-destination-mac`	
Related Commands		
Note		

show interfaces port-channel summary

show interfaces port-channel summary

Displays a summary for LAG interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.4.1100 Updated Example
Role	admin

Example	`switch (config) # show interfaces port-channel summary`
	`Flags: D - Down, U - Up, P - Up in port-channel (members)`
	`S - Suspend in port-channel (members), I - Individual`
	
	`----------------------------------------------------------------------`
	`Group Port- Type Member Ports`
	`Channel`
	`----------------------------------------------------------------------`
	`1 Po2(U) LACP Eth1/58(D) Eth1/59(I) Eth1/60(S)`
	`2 Po5(D) LACP Eth1/1(S) Eth1/33(I)`
	`3 Po10(U) LACP Eth1/49(P) Eth1/50(P) Eth1/51(S) Eth1/52(S)`

Related Commands	
Note	

5.4 Link Layer Discovery Protocol (LLDP)

The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.

5.4.1 Configuring LLDP

➢ ***To configure the LLDP on the switch:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable LLDP globally on the switch. Run:

```
switch (config) # lldp
switch (config) #
```

Step 4. Enable LLDP per interface. Run:

```
switch (config  1/1) # lldp receive
switch (config  1/1) # lldp transmit
```

Step 5. Display LLDP local information. Run:

```
switch (config) # show lldp local

LLDP is Enabled

Local global configuration
Chassis sub type: macAddress (4)
Chassis id: 00:11:22:33:44:55
System Name: "switch-111111"
System Description: my-system-description
Supported capabilities: B
Supported capabilities enabled: B
```

Step 6. Display LLDP remote information. Run:

```
switch (config)# show lldp interfaces ethernet 1/1 remote

Ethernet 1/1
Remote Index: 1
Remote chassis id: 00:11:22:33:44:55 ; chassis id subtype: mac
Remote port-id: ethenret 1/2; port id subtype: local
Remote port description: ethernet 1/2
Remote system name: remote-system
Remote system description: remote-system-description
Remote system capabilities supported: B ; B
```

5.4.2 DCBX

Data Center Bridging (DCB) is an enabler for running the Ethernet network with lossless connectivity using priority-based flow control and enhanced transmission selection. DCBX (exchange) complements the DCB implementation by offering a dynamic protocol that communicates DCB attributes between peering endpoint.

Onyx supports two versions of DCBX TLVs running on top of LLDP:

- DCBX IEEE
- DCBX CEE

By default DCBX IEEE is enabled when LLDP is enabled (LLDP, however, is not enabled by default).

For more information, please refer to the Mellanox Community at:
https://community.mellanox.com/docs/DOC-2485.

5.4.3 Commands

lldp

	lldp **no lldp** Enables LLDP globally. The no form of the command disables the LLDP.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.2.0300
Role	admin
Example	`switch (config)# lldp`
Related Commands	show lldp local
Note	

lldp reinit

lldp reinit <seconds>
no lldp reinit

Sets the delay in seconds from enabling the LLDP on the port until re-initialization will be attempted.
The no form of the command sets the parameter to default.

Syntax Description	seconds	1-10
Default	2	
Configuration Mode	config	
History	3.2.0300	
Role	admin	
Example	`switch (config)# lldp reinit 10`	
Related Commands	show lldp timers	
Note		

lldp timer

lldp timer <seconds>
no lldp timer

Sets the LLDP interval at which LLDP frames are transmitted. (lldpMessageTxInterval)
The no form of the command sets the parameter to default.

Syntax Description	seconds	5-32768
Default	30	
Configuration Mode	config	
History	3.2.0300	
Role	admin	
Example	`switch (config)# lldp timer 10`	
Related Commands	show lldp timers	
Note		

lldp tx-delay

lldp tx-delay <seconds>
no lldp tx-delay

Indicates the delay in seconds between successive LLDP frame transmissions
The no form of the command sets the parameter to default.

Syntax Description	seconds 1-8192
Default	2
Configuration Mode	config
History	3.2.0300
Role	admin
Example	`switch (config)# lldp tx-delay 10`
Related Commands	show lldp timers
Note	The recommended value for the tx-delay is set by the following formula: 1 <= lldp tx-delay <= (0.25 * lldp timer)

lldp tx-hold-multiplier

lldp tx-hold-multiplier <seconds>
no lldp tx-hold-multiplier

The time-to-live value expressed as a multiple of the lldpMessageTxInterval object. The no form of the command sets the parameter to default.

Syntax Description	seconds 1-8192
Default	2
Configuration Mode	config
History	3.2.0300
Role	admin
Example	`switch (config)# lldp tx-hold-multiplier 10`
Related Commands	show lldp timers
Note	The actual time-to-live value used in LLDP frames, can be expressed by the following formula: TTL = min(65535, (lldpMessageTxInterval * lldpMessageTxHoldMultiplier)) For example, if the value of lldpMessageTxInterval is '30', and the value of lldpMessageTxHoldMultiplier is '4', then the value '120' is encoded in the TTL field in the LLDP header.

lldp (interface)

lldp {receive | transmit}
no lldp {receive | transmit}

Enables LLDP receive or transmit capabilities.
The no form of the command disables LLDP receive or transmit capabilities.

Syntax Description	med-tlv-select	Enables LLDP media TLVs
	receive	Enables LLDP receive on this port
	tlv-select	Enables LLDP TLVs
	transmit	Enables LLDP transmit on this port
Default	Enabled for receive and trasmit.	
Configuration Mode	config interface ethernet	
History	3.2.0300	
Role	admin	
Example	`switch (config  1/1)# lldp receive`	
Related Commands	show lldp interface	
Note	The LLDP is disabled by default (globally)	

lldp tlv-select

lldp tlv-select {[dcbx] [dcbx-cee] [port-description] [sys-name] [sys-description] [sys-capababilities] [management-address] [none] all}

Sets the LLDP basic TLVs to be transmitted on this port.

Syntax Description	dcbx	Enables LLDP-DCBX TLVs
	dcbx-cee	Enables LLDP-DCBX CEE TLVs
	port-description	LLDP port description TLV
	sys-name	LLDP system name TLV
	sys-description	LLDP system description TLV
	sys-capabilities	LLDP system capabilities TLV
	management-address	LLDP management address TLV
	all	all above TLVs
	none	None of the above TLVs
Default	all	
Configuration Mode	config interface ethernet	
History	3.2.0300	
	3.3.0000	Added “none” parameter
	3.3.4302	Added “dcbx” parameter
	3.3.4402	Added “dcbx-cee” parameter
Role	admin	
Example	`switch (config  1/1)# lldp tlv-select port-description sys-name`	
Related Commands	show lldp interface	
Note	The management address is chosen according to the following criteria where 1 takes priority over 2, and 2 takes priority over 3: 1. Smallest IP address of mgmt0 2. Smallest IP address of mgmt1 3. First primary address of all non-management interfaces	

lldp med-tlv-select

lldp med-tlv-select {all | media-capability | network-policy | none}

Configures LLDP media TLV attributes.

Syntax Description	all	Enables all LLDP media TLVs
	media-capabilities	Enables Media Capabilities TLV
	network-policy	Enables Network-Policy TLV
	none	Disables all LLDP media TLVs
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1)# lldp med-tlv-select all`	
Related Commands	show lldp interface	
Note		

dcb application-priority

dcb application-priority <selector> <protocol> <priority>

Adds an application to the application priority table.

Syntax Description	selector	Protocol type: ethertype
	protocol	Protocol field in hexadecimal notation (e.g. '0x8906' for FCoE, '0x8914' for FIP).
	priority	Range: 0-7.
Default	No applications are available. The table is empty.	
Configuration Mode	config	
History	3.3.4200	
	3.4.0008	
Role	admin	
Example	`switch (config-if)# dcb application-priority ethertype 0x8906`	
Related Commands	show lldp interface	
Note		

clear lldp counters

clear lldp counters [<Device | Port>]

Clears LLDP counters for all ports or for a specific port.

Syntax Description	N/A
Default	N/A
Configuration Mode	Global
History	3.6.4006
Role	admin
Example	`switch (config) # clear lldp counters` `switch (config) # clear lldp counters  1/1`
Related Commands	N/A
Note	

show lldp local

	show lldp local Displays LLDP local information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300
Role	admin

Example	`switch (config)# show lldp local`
	
	`LLDP is Enabled`
	
	`Local global configuration`
	
	`Chassis sub type: macAddress (4)`
	`Chassis id: 0002C9030046AF00`
	`System Name: my-switch`
	`System Description: SN2100`
	`Supported capabilities: B,R`
	`Supported capabilities enabled: B`

Related Commands	
Note	

show lldp interfaces

show lldp interfaces [ethernet <inf> [med-cap | remote]]

Displays LLDP remote interface table information.

Syntax Description	inf	Local interface number (e.g. 1/1)
	med-cap	Displays local port media capabilities information
	remote	Displays LLDP Ethernet remote configuration & status
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0300	
	3.3.4200	Updated Example
	3.6.1002	Updated Example
Role	admin	
Example		

```
switch (config)# show lldp interfaces
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA:
management-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC:
Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------
Eth1/1    Enabled   Enabled   PD, SD
Eth1/2    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/3    Disabled  Disabled  PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-NWP
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R,
                              MED-CAP, MED-NWP
Eth1/5    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/6    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/7    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
```

Related Commands	
Note	

show lldp remote

	show lldp remote Displays LLDP remote information (remote device id, remote port id, remote system name).
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Role	admin

Example	`switch (config)# show lldp remote`
	`---------------------------------------------------------------------------`
	`Local Interface  Device ID  Port ID  System Name`
	`---------------------------------------------------------------------------`
	`Eth1/4  e4:1d:2d:a5:f3:35  e4:1d:2d:a5:f3:35  Not Advertised`
	`Eth1/10  e4:1d:2d:44:65:00  Eth1/10  arc-switch108`
	`Eth1/11  e4:1d:2d:44:65:00  Eth1/11  arc-switch108`
	`Eth1/12  e4:1d:2d:44:65:00  Eth1/12  arc-switch108`
	`Eth1/13  e4:1d:2d:44:65:00  Eth1/13  arc-switch108`
	`Eth1/14  e4:1d:2d:44:65:00  Eth1/14  arc-switch108`
	`Eth1/15  e4:1d:2d:44:65:00  Eth1/15  arc-switch108`
	`Eth1/16  e4:1d:2d:44:65:00  Eth1/16  arc-switch108`
	`Eth1/17  e4:1d:2d:44:65:00  Eth1/17  arc-switch108`
	`Eth1/18  e4:1d:2d:44:65:00  Eth1/18  arc-switch108`
	`Eth1/19  e4:1d:2d:44:65:00  Eth1/19  arc-switch108`
	`Eth1/20  e4:1d:2d:44:65:00  Eth1/20  arc-switch108`
	`Eth1/21  e4:1d:2d:44:65:00  Eth1/21  arc-switch108`
	`Eth1/22  e4:1d:2d:44:65:00  Eth1/22  arc-switch108`
	`Eth1/23  e4:1d:2d:44:65:00  Eth1/23  arc-switch108`
	`Eth1/24  e4:1d:2d:44:65:00  Eth1/24  arc-switch108`
	`Eth1/25  e4:1d:2d:44:65:00  Eth1/25  arc-switch108`
	`Eth1/26  e4:1d:2d:44:65:00  Eth1/26  arc-switch108`
	`Eth1/31  e4:1d:2d:44:65:00  Eth1/31  arc-switch108`
	`Eth1/32  e4:1d:2d:44:65:00  Eth1/32  arc-switch108`

Related Commands	
Note	

show lldp statistics [<inf>]

	show lldp statistics [<inf>] Displays LLDP interface statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300
Role	admin

Example	`switch (config)# show lldp statistics  1/1`
	`Interface Frames    In     In    TLVs      TLVs        Ageout Out`
	`          Discarded Errors Total Discarded Unrecognize        Frames`
	`---------------------------------------------------------------------`
	`Eth 1/1       0        0     10    0         0           0      0`
	`switch (config)#`

Related Commands	
Note	

show lldp statistics global

	show lldp statistics global Displays LLDP global statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300
Role	admin

Example	`switch (config)# show lldp timers`
	`Remote Table Last Change Time : 10300`
	`Remote Table Inserts : 5`
	`Remote Table Deletes : 0`
	`Remote Table Drops : 0`
	`Remote Table Ageouts : 0`
	`switch (config)#`

Related Commands	
Note	

show lldp timers

	show lldp timers Displays LLDP timers configuration
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300
Role	admin

Example	`switch (config)# show lldp timers`
	`msg-tx-interval:30`
	`tx-delay:2`
	`tx-hold:4`
	`tx-reinit-delay:2`
	`switch (config)#`

Related Commands	
Note	

show dcb application-priority

	show dcb application-priority Displays application priority admin table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4200
Role	admin

Example	`switch (config)# show dcb application-priority`
	
	`Application priority configuration`
	
	`Selector  Protocol  Priority`
	`----------------------------`
	`Ethertype  0x8906  3`
	`Ethertype  0x8914  3`
	
	`switch (config)#`

Related Commands	
Note	

5.5 VLANs

A Virtual Local Area Network (VLAN) is an L2 segment of the network which defines a broadcast domain and is identified by a tag added to all Ethernet frames running within the domain. This tag is called a VLAN ID (VID) and can take a value of 1-4094.

Each port can have a switch mode of either:

- Access – access port is a port connected to a host. It can accept only untagged frames, and assigns them a default configured VLAN (Port VLAN ID). On egress, traffic sent from the access port is untagged.
- Access-dcb – this mode is Mellanox specific that receives ingress untagged traffic but sends egress priority tag (VLAN ID = 0)
- Hybrid – hybrid port is a port connected to either switches or hosts. It can receive both tagged and untagged frames and assigns untagged frames a default configured VLAN (Port VLAN ID). It receives tagged frames with VLANs of which the port is a member (these VLANs' names are allowed). On egress, traffic of allowed VLANs sent from the Hybrid port is sent tagged, while traffic sent with PVID is untagged.
- Trunk – trunk port is a port connecting 2 switches. It accepts only tagged frames with VLANs of which the port is a member. On egress, traffic sent from the Trunk port is tagged. By default, a Trunk port is, automatically, a member on all current VLANs.

5.5.1 Configuring Access Mode and Assigning Port VLAN ID (PVID)

➢ ***To configure Access mode and assign PVID to interfaces:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 6
switch (config vlan 6) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 6) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) #
```

Step 6. From within the interface context, configure the interface mode to Access. Run:

```
switch (config interface ethernet 1/22) # switchport mode access
```

Step 7. From within the interface context, configure the Access VLAN membership. Run:

```
switch (config interface ethernet 1/22) # switchport access vlan 6
```

Step 8. Change back to config mode. Run:

```
switch (config  1/22) # exit
switch (config) #
```

5.5.2 Configuring Hybrid Mode and Assigning Port VLAN ID (PVID)

➢ ***To configure Hybrid mode and assign PVID to interfaces:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 6
switch (config vlan 6) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 6) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) #
```

Step 6. From within the interface context, configure the interface mode to Access. Run:

```
switch (config interface ethernet 1/22) # switchport mode hybrid
switch (config interface ethernet 1/22) #
```

Step 7. From within the interface context, configure the Access VLAN membership. Run:

```
switch (config interface ethernet 1/22) # switchport access vlan 6
switch (config interface ethernet 1/22) #
```

Step 8. Change to config mode again. Run:

```
switch (config interface ethernet 1/22) # exit
switch (config) #
```

5.5.3 Configuring Trunk Mode VLAN Membership

➢ ***To configure Trunk mode VLAN membership:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 10
switch (config vlan 10) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 10) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch [standalone: master] (config) # interface ethernet 1/35
switch [standalone: master] (config interface ethernet 1/35) #
```

Step 6. From within the interface context, configure the interface mode to Trunk. Run:

```
switch [standalone: master] (config interface ethernet 1/35) # switchport mode trunk
switch [standalone: master] (config interface ethernet 1/35) #
```

5.5.4 Configuring Hybrid Mode VLAN Membership

➢ ***To configure Hybrid mode VLAN membership:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 10
switch (config vlan 10) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 10) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/35
switch (config interface ethernet 1/35) #
```

Step 6. From within the interface context, configure the interface mode to Hybrid. Run:

```
switch (config interface ethernet 1/35) # switchport mode hybrid
switch (config interface ethernet 1/35) #
```

Step 7. From within the interface context, configure the allowed VLAN membership. Run:

```
switch (config interface ethernet 1/35) # switchport hybrid allowed-vlan add 10
switch (config interface ethernet 1/35) #
```

Step 8. Change to config mode again. Run:

```
switch (config interface ethernet 1/35) # exit
switch (config) #
```

5.5.5 Commands

vlan

vlan {<vlan-id> | <vlan-range>}
no vlan {<vlan-id> | <vlan-range>}

Creates a VLAN or range of VLANs, and enters a VLAN context.
The no form of the command deletes the VLAN or VLAN range.

Syntax Description	vlan-id	Range: 1-4094
	vlan-range	Any range of VLANs
Default	VLAN 1 is enabled by default.	
Configuration Mode	config	
History	3.1.1400	
Role	admin	
Example	`switch (config) # vlan 10` `switch (config vlan 10) #`	
Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan	
Note	Interfaces are not added automatically to VLAN unless configured with trunk or hybrid mode with “all” option turned on.	

name

name <vlan-name>
no name

Adds VLAN name.
The no form of the command deletes the VLAN name.

Syntax Description	vlan-name	40-character long string
Default	No name available	
Configuration Mode	config vlan	
History	3.1.1400	
Role	admin	
Example	`switch (config vlan 10) # name my-vlan-name`	
Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan	
Note	Name can not be configured for a range of VLANs.	

show vlan

show vlan [id <vlan-id>]

Displays the VLAN table.

Syntax Description	vlan-id 1-4094.
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400
Role	admin

Example	`switch (config vlan 10) # show vlan`
	
	`VLAN Name Ports`
	`---- ---------- ----------`
	`1 default Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...`
	`10 my-vlan-name`

Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan vlan
Note	

switchport mode

switchport mode {access | dot1q-tunnel | trunk | hybrid | access-dcb}
no switchport mode

Sets the switch port mode.
The no form of the command sets the switch port mode to access.

Syntax Description	access	Untagged port. 802.1q tagged traffic are filtered. Egress traffic is untagged.
	dot1q-tunnel	Allows both tagged and untagged ingress Ethernet packets. Egress packets are tagged with a second VLAN (802.1Q) header.
	trunk	802.1q tagged port, untagged traffic is filtered.
	hybrid	Both 802.1q tagged and untagged traffic is allowed on the port.
	access-dcb	Untagged port, egress traffic is priority tagged.
Default	access	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
	3.3.4500	Added MPO configuration mode
	3.4.3000	Added dot1q-tunnel parameter
	3.6.6000	Added ability to switchport mode for a range of interfaces
Role	admin	
Example	`switch (config) # interface ethernet 1/7` `switch (config interface ethernet 1/7) # switchport mode access`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan switchport dot1q-tunnel qos-mode vlan	
Note	Switchport mode may be configured for a range of interfaces (interface <inf-type> <id-range> switchport mode <type>)	

switchport dot1q-tunnel qos-mode

switchport dot1q-tunnel qos-mode {pipe | uniform}
no switchport dot1q-tunnel qos-mode

Assigns QoS to the service provider's traffic.
The no form of the command resets the parameter value to its default.

Syntax Description	pipe	Gives the service provider's traffic QoS 0
	uniform	Gives the service provider's traffic the same QoS as the customer's traffic
Default	pipe	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.4.3000	
Role	admin	
Example	`switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos-mode uniform` `switch (config interface ethernet 1/1) #`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan vlan	
Note		

switchport access

	switchport access vlan <vlan-id> **no switchport access vlan** **switchport access none** (hybrid mode only) Sets the port access VLAN. The no form of the command sets the port access VLAN to 1. The none clause of the command removes access vlan membership from the port, thus disallowing untagged traffic on this port. This is commonly used for fast transaction from hybrid switchport to trunk-like switchport and vice versa.
Syntax Description	vlan-id 1-4094.
Default	1
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.1.1400 3.2.0500 Format change (removed hybrid and access-dcb options). Previous command format was: “switchport {hybrid | access-dcb | access} vlan <vlan-id>” 3.3.4500 Added MPO configuration mode 3.6.6000 Added ability to configure VLAN ID for a range of interfaces 3.7.11xx Updated command syntax & notes.
Role	admin
Example	`switch (config) #  1/7` `switch (config  1/7) # switchport access vlan 10`
Related Commands	show vlan show interfaces switchport switchport mode switchport [trunk | hybrid] allowed-vlan vlan
Note	• This command is not applicable for interfaces with port mode trunk. • Only one option (“access”, “access-dcb” or “hybrid”) is possible to configure on the port, depending on the switchport mode of the port. • Access VLAN ID may be configured to a range of interfaces (interface <inf-type> <id-range> switchport access vlan <vlan-ID>) • This command is not applicable for interfaces with port mode trunk. • In hybrid mode, access vlan is optional. Alternatively, use “access none” in order to disable access vlan. In this case, all incoming untagged traffic will be dropped.

switchport {hybrid, trunk} allowed-vlan

switchport {hybrid, trunk} allowed-vlan {<vlan> | add <vlan> | remove <vlan> all | except <vlan> | none}

Sets the port allowed VLANs.

Syntax Description	vlan	VLAN ID (1-4094) or VLAN range
	add	Adds VLAN or range of VLANs
	remove	Removes VLANs or range of VLANs
	all	Adds all VLANs in available in the VLAN table. New VLANs added to the VLAN table are added automatically.
	except	Adds all VLANs expect this VLAN or VLAN range
	none	Removes all VLANs
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
Role	admin	
Example	`switch (config interface ethernet 1/7) # switchport hybrid allowed-vlan all`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport mode vlan	
Note	• This command is not applicable for interfaces with port mode access or access-dcb • In order for the parameter “hybrid” or “trunk” to be available, the switchport mode on the interface must be configured to either hybrid or trunk respectively	

switchport voice

switchport voice vlan <vlan-id>
no switchport voice vlan

Configures voice VLAN for the interface.
The no form of the command disables voice VLAN.

Syntax Description	vlan-id 1-4094.
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config) # interface ethernet 1/7` `switch (config interface ethernet 1/7) # switchport voice vlan 10`
Related Commands	lldp med-tlv-select show vlan show interfaces switchport switchport mode switchport [trunk | hybrid] allowed-vlan vlan
Note	

show interfaces switchport

show interfaces [<if>] switchport

Displays all interface switch port configurations.

Syntax Description	if	Possible interface types: • ethernet <slot/port> • port-channel <lag-id> • mlag-port-channel <id>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.6102	Added ability to filter by specific interfaces and updated Example
Role	admin	

Example

```
switch (config) # show interfaces switchport

-------------------------------------------------------------------
Interface      Mode        Access vlan       Allowed vlans
-------------------------------------------------------------------
Eth1/1           trunk          N/A                  3-10
Eth1/2           access         1
Eth1/3           access         1
Eth1/4           access         1
Eth1/5           access         1
Eth1/6           access         1
Eth1/7           access         1
Eth1/8           access         1
Eth1/9           access         1
Eth1/10          access         1
Eth1/11          access         1
Eth1/12          access         1
Eth1/13          access         1
Eth1/14          access         1
Eth1/15          access         1
Eth1/16          access         1
Po1              hybrid         5                    6-8
Mpo2             trunk          N/A                  1-3, 5-10
```

Related Commands	show vlan switchport access vlan switchport mode vlan
Note	This command can accept an explicit interface or interface range (displays information only for available interfaces)

5.6 Spanning Tree

The operation of Rapid Spanning Tree Protocol (RSTP) provides for rapid recovery of connectivity following the failure of a bridge/bridge port or a LAN. The RSTP component avoids this delay by calculating an alternate root port, and immediately switching over to the alternate port if the root port becomes unavailable. Thus, using RSTP, the switch immediately brings the alternate port to forwarding state, without the delays caused by the listening and learning states. The RSTP component conforms to IEEE standard 802.1D 2004.

RSTP enhancements is a set of functions added to increase the volume of RSTP in Mellanox switches. It adds a set of capabilities related to the behavior of ports in different segments of the network. For example: the required behavior of a port connected to a non-switch entity, such as host, is to converge quickly, while the required behavior of a port connected to a switch entity is to converge based on the RSTP parameters.

Additionally, it adds security issues on a port and switch basis, allowing the operator to determine the state and role of a port or the entire switch should an abnormal event occur. For example: If a port is configured to be root-guard, the operator will not allow it to become a root-port under any circumstances, regardless of any BPDU that will have been received on the port.

5.6.1 Port Priority and Cost

When two ports on a switch are part of a loop, the STP port priority and port path cost configuration determine which port on the switch is put in the forwarding state and which port is put in the blocking state.

To configure port priority use the following command:

```
switch (config  <inf>)# spanning-tree port-priority <0-240>
```

To configure port path cost use the following command:

```
switch (config  <inf>)# spanning-tree cost <1-200000000>
```

5.6.2 Port Type

Port type has the following configuration options:

- **edge –** is not assumed to be converged by the RSTP learning/forwarding mechanism. It converges to forwarding quickly.

It is recommended to configure the port type for all ports connected to hosts as edge ports.

- **normal –** is assumed to be connected to a switch, thus it tries to be converged by the RSTP learning/forwarding. However, if it does not receive any BPDUs, it is operationally moved to be edge.
- **network –** is assumed to be connected only to a switch or bridge.

Each of these configuration options is mutually exclusive.

Port type is configured using the command spanning-tree port type. It may be applied globally on the switch (Config) level, which configures all switch interfaces. Another option is to configure ports individually by entering the interface's configuration mode.

- Global configuration:

```
switch (config)# spanning-tree port type {edge , normal , network} default
```

- Interface configuration:

```
switch (config  <inf>)# spanning-tree port type {edge , normal, network}
```

5.6.3 BPDU Filter

Using BPDU filter prevents the CPU from sending/receiving BPDUs on specific ports.

BPDU filtering is configured per interface. When configured, the port does not send any BPDUs and drops all BPDUs that it receives. To configure BPDU filter, use the following command:

```
switch (config  <inf>)# spanning-tree bpdufilter {enable , disable}
```

Configuring BPDU filtering on a port connected to a switch can cause bridging loops because the port filters any BPDU it receives and goes to forwarding state.

5.6.4 BPDU Guard

BPDU guard is a security feature which, when enabled, shuts down the port in case it receives BPDU packets. This feature becomes useful when connecting to an unauthorized switch.

To configure BPDU guard use the following command:

```
switch (config  <inf>)# spanning-tree bpduguard {enable , disable}
```

5.6.5 Loop Guard

Loop guard is a feature that prevents loops in the network.

When a blocking port in a redundant topology transitions to the forwarding state (accidentally), an STP loop occurs. This happens when BPDUs are no longer received by one of the ports in a physically redundant topology.

Loop guard is useful in switched networks where devices are connected point-to-point. A designated bridge cannot disappear unless it sends an inferior BPDU or brings the link down on a point-to-point connection.

The loop guard configuration is only allowed on "network" and "normal" port types.

If loop guard is enabled and the port does not receive BPDUs, the port is put into an inconsistent state (blocking) until the port starts to receive BPDUs again. A port in the inconsistent state does not transmit BPDUs. If BPDUs are received again, loop guard alters its inconsistent state condi-

tion. STP converges to a stable topology without the failed link or bridge after loop guard isolates the failure.

Disabling loop guard moves all loop-inconsistent ports to listening state.

To configure loop guard use the following command:

```
switch (config  <inf>)# spanning-tree guard loop
```

5.6.6 Root Guard

Configuring root guard on a port prevents that port from becoming a root port. A port put in root-inconsistent (blocked) state if an STP convergence is triggered by a BPDU that makes that port a root port. The port is unblocked after the port stops sending BPDUs.

To configure loop guard use the following command:

```
switch (config  <inf>)# spanning-tree guard root
```

5.6.7 MSTP

Spanning Tree Protocol (STP) is a mandatory protocol to run on L2 Ethernet networks to eliminate network loops and the resulting broadcast storm caused by these loops. Multiple STP (MSTP) enables the virtualization of the L2 domain into several VLANs, each governed by a separate instance of a spanning tree which results in a network with higher utilization of physical links while still keeping the loop free topology on a logical level.

Up to 64 MSTP instances can be configured on a switch.

For MSTP network design over Mellanox L2 VMS, please refer to Mellanox Virtual Modular Switch Reference Guide.

5.6.8 RPVST

Rapid Per-VLAN Spanning Tree (RPVST) flavor of the STP provides finer-grained traffic by paving a spanning-tree instance per each configured VLAN. Like MSTP, it allows a better utilization of the network links comparing to RSTP.

Figure 17 exhibits a typical RPVST network configuration to get a better utilization on the inter-switch trunk ports.

Figure 17: RPVST Network Config

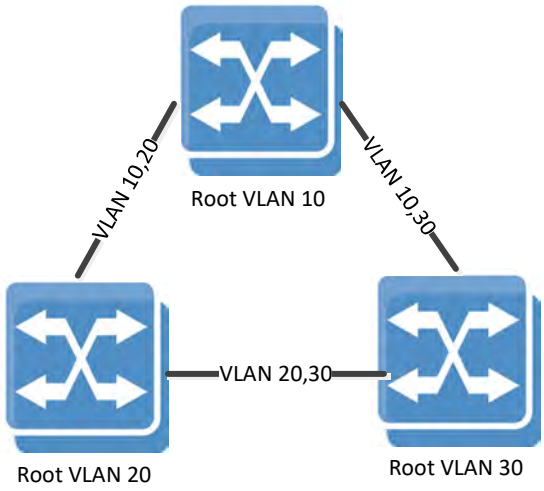

5.6.8.1 RPVST and VLAN Limitations

When the STP of the switch is set to RPVST, spanning tree is set on each of the configured VLANs in the system by default. To enable the spanning tree mode, the command "spanning-tree" must be run.

Each VLAN runs an STP state machine and an RPVST instance. There is a global limitation on the number of active state machines that can operate in Onyx. Enforcement of this limitation is done through the maximum number of VLANs allowed in the system (128).

The state machine takes attributes like forward time, hello time, max age and priority, etc.

When configuring priority on a VLAN in RPVST, the operational priority given to the VLAN is a summation of what the user configured and the value of the VLAN itself. For example running "spanning-tree vlan 10 priority 32768" yields a priority of 32778 for VLAN 10.

5.6.8.2 RPVST and RSTP Interoperability

Figure 18: RPVST and RSTP Cluster

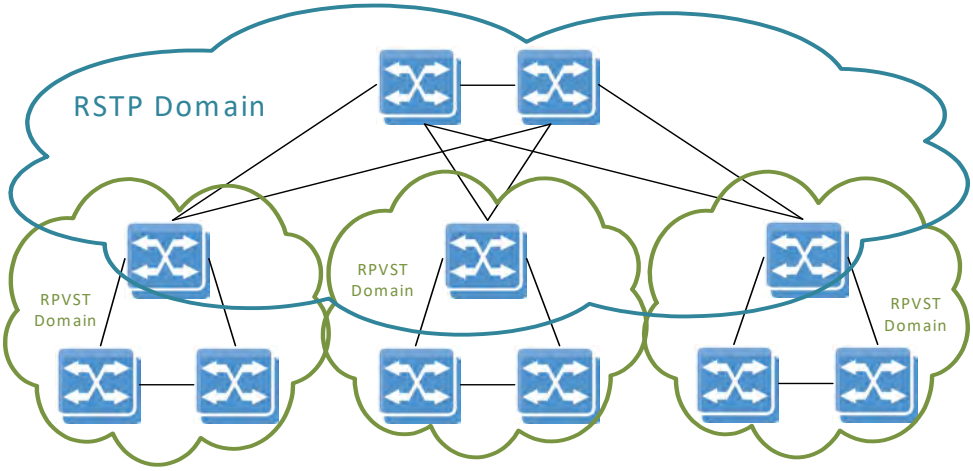

RPVST domains can be interconnected by a standard 802.1Q domain that runs RSTP protocol. While the RSTP domain builds a single common instance spanning tree, the RPVST domains at the edge continue to build a tree per VLAN while exchanging tagged RPVST multicast BPDUs.

(This exchange may happen on untagged RPVST BPDUs as well.) The switch devices that are in the boundary between the RPVST and the RSTP domains should be configured as RPVST mode.

When set to RPVST mode, the switch continues to run the common instance spanning tree (CIST) state machine on VLAN 1 by exchanging IEEE BPDUs with the legacy RSTP switches.

To successfully connect RSTP and RPVST domains, the system administrator must align the native VLAN configuration across all network switches, or in other words, the internal identification of untagged packets to VLAN.

5.6.9 Commands

spanning-tree

spanning-tree
no spanning-tree

Globally enables the spanning tree feature.
The no form disables the spanning tree feature.

Syntax Description	N/A
Default	Spanning tree is enabled.
Configuration Mode	config
History	3.1.0000
Role	admin
Example	switch (config) # no spanning-tree
Related Commands	show spanning-tree
Note	

spanning-tree mode

spanning-tree mode {rst | mst | rpvst}
no spanning-tree mode

Changes the spanning tree mode.
The no form of the command sets the parameter to its default value.

Syntax Description	mst	Multiple spanning tree.
	rst	Rapid spanning tree.
	rpvst	Rapid per-VLAN spanning tree.
Default	rst	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mode mst`	
Related Commands		
Note	The number of VLANs supported by RPVST is 128	

spanning-tree (timers)

spanning-tree [forward-time <time in secs> | hello-time <time in secs> | max-age <time in secs>]
no spanning-tree [forward-time | hello-time | max-age | priority]

Sets the spanning tree timers.
The no form of the command sets the timer to default.

Syntax Description	forward-time	Controls how fast a port changes its spanning tree state from Blocking state to Forwarding state. Parameter range: 4-30 seconds.
	hello-time	Determines how often the switch broadcasts its hello message to other switches when it is the root of the spanning tree. Parameter range: 1-2 seconds.
	max-age	Sets the maximum age allowed for the Spanning Tree Protocol information learnt from the network on any port before it is discarded. Parameter range: 6-40 seconds.
Default	forward-time: 15 seconds hello-time:2 seconds max-age: 20 seconds	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # spanning-tree forward-time` `switch (config) #`	
Related Commands	show spanning-tree	
Note	The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1)	

spanning-tree port type (default global)

	spanning-tree port type {edge [bpdufilter | bpduguard] | network [bpduguard] | normal [bpduguard]} default **no spanning-tree port type default** Configures all switch interfaces as edge/network/normal ports. These ports can be connected to any type of device. The no form of the command disables the spanning tree operation.	
Syntax Description	edge	Assumes all ports are connected to hosts/servers.
	bpdufilter	Configures to enable the spanning tree BPDU filter.
	bpduguard	Configures to enable the spanning tree BPDU guard.
	network	Assumes all ports are connected to switches and bridges.
	normal	The port type (edge or network) determines according to the spanning tree operational mode.
Default	Normal	
Configuration Mode	config	
History	3.1.0000	
	3.4.0008	Updated command syntax
Role	admin	
Example	`switch (config) # spanning-tree port type edge default` `switch (config) #`	
Related Commands	show spanning-tree	
Note		

spanning-tree priority

	spanning-tree priority <bridge-priority> **no spanning-tree priority** Sets the spanning tree bridge priority. The no form of the command sets the bridge priority to default.	
Syntax Description	bridge-priority	Sets the bridge priority for the spanning tree. Its value must be in steps of 4096, starting from 0. Only the following values are applicable: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440.
Default	32786	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	`switch (config) # spanning-tree priority 4096` `switch (config) #`	
Related Commands	show spanning-tree	
Note		

spanning-tree port-priority

spanning-tree port-priority <priority>
no spanning-tree port-priority

Configures the spanning-tree interface priority.
The no form of the command returns configuration to its default.

Syntax Description	priority	Spanning tree interface priority. The possible values are: 0, 16, 32,48, 64, 80, 96, 112, 128,144, 160, 176, 192, 208, 224, 240.
Default	128	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	

Example	`switch (config) #  1/1`
	`switch (config  1/1) # spanning-tree port-priority 16`
	`switch (config  1/1) #`	

Related Commands	show spanning-tree	
Note		

spanning-tree cost

	spanning-tree cost <port cost> **no spanning-tree cost** Configures the interface cost of the spanning tree. The no form of the command returns configuration to its default.	
Syntax Description	port cost	Sets the spanning tree cost of an interface. Value range is 0-200000000.
Default	The default cost is derived from the speed. 1Gbps 20000 10Gbps 2000 40Gbps 500 50Gbps 400 100Gbpg 200	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	

Example	`switch (config) #  1/1`
	`switch (config  1/1) # spanning-tree cost 1000`
	`switch (config  1/1) #`	

Related Commands	show spanning-tree	
Note	• LAG default cost is calculated by dividing the port speed by the number of active links in UP state. For example: if there were 4 links in the LAG out of which only two are in UP state, assuming the port speed is 10Gbps, the LAG cost will be 2000/2 = 1000. • When configuring the cost for a LAG, the cost will be fixed to this configuration, no matter what the number of active links (UP state) in the LAG is • Unstable network may cause the LAG cost to change dynamically assuming the cost parameter is not configured for anything else other than default	

spanning-tree port type

	spanning-tree port type <port type> **no spanning-tree port type** Configures spanning-tree port type The no form of the command returns configuration to default.	
Syntax Description	default	According to global configuration
	edge	Assumes all ports are connected to hosts/servers.
	normal	The port type (edge or network) determines according to the spanning tree operational mode.
	network	Assumes all ports are connected to switches and bridges.
	bpdufilter	Configures to enable the spanning tree BPDU filter.
	bpduguard	Configures to enable the spanning tree BPDU guard.
Default	Globally defined by the command “spanning-tree port type <port-type> default”	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	

Example	`switch (config) #  1/1`
	`switch (config  1/1) # spanning-tree port type edge`
	`switch (config  1/1) #`	

Related Commands	show spanning-tree	
Note		

spanning-tree guard

	spanning-tree guard {loop | root} **no spanning-tree guard {loop | root}** Configures spanning-tree guard. The no form of the command returns configuration to default.	
Syntax Description	loop	Enables loop-guard on the interface. If the loop-guard is enabled, upon a situation where the interface fails to receive BPDUs the switch will not egress data traffic on this interface.
	root	Enables root-guard on the interface. If root-guard is enabled on the interface, the interface will never be selected as root port.
Default	loop-guard and root-guard are disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Role	admin	

Example	`switch (config) #  1/1`
	`switch (config  1/1) # spanning-tree guard root`
	`switch (config  1/1) #`	

Related Commands	show spanning-tree	
Note		

spanning-tree bpdufilter

spanning-tree bpdufilter {disable | enable}
no spanning-tree bpdufilter

Configures spanning-tree BPDU filter on the interface. The interface will ignore any BPDU that it receives and will not send PDBUs, The STP state on the port will move to the forwarding state.
The no form of the command returns the configuration to default.

Syntax Description	disable	Disables the BPDU filter on this port.
	enable	Enables the BPDU filter on this port.
Default	BPDU filter is disabled.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
Role	admin	
Example	`switch (config) # interface ethernet 1/1` `switch (config interface ethernet 1/1) # spanning-tree bpdufilter enable`	
Related Commands	show spanning-tree	
Note	This command can be used when the switch is connected to hosts.	

clear spanning-tree counters

	clear spanning-tree counters Clears the spanning-tree counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config) # clear panning-tree counters`
Related Commands	show spanning tree
Note	

spanning-tree mst max-hops

spanning-tree mst max-hops <max-hops>
no spanning-tree mst max-hops

Specifies the max hop value inserts into BPDUs that sent out as the root bridge. The no form of the command sets the parameter to its default value.

Syntax Description	max-hops	Max hop value. The range is 6-40.
Default	20	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mst max-hops 20`	
Related Commands		
Note	• The max hop setting determines the number of bridges in an MST region that a BPDU can traverse before it is discarded • This command is available when global STP mode is set to MST	

spanning-tree mst priority

spanning-tree mst <mst-instance> priority <priority>
no spanning-tree mst <mst-instance> priority

Configures the specified instance's priority number.
The no form of the command sets the parameter to its default value.

Syntax Description	mst-instance	MST instance. Range is 1-64.
	priority	MST instance port priority. Possible values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440
Default	32768	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mst 1 priority 32768`	
Related Commands		
Note	• The bridge priority is the four most significant digits of the bridge ID, which is used by spanning tree algorithms to select the root bridge and choose among redundant links. Bridge ID numbers range from 0-65535 (16 bits); bridges with smaller bridge IDs are elected over other bridges. • This command is available when global STP mode is set to MST	

spanning-tree mst vlan

	spanning-tree mst <mst-instance> vlan <vlan-range> **no spanning-tree mst <mst-instance> vlan <vlan-range>** Maps a VLAN or a range of VLANs into an MSTP instance. The no form of the command unmaps a VLAN or a range of VLANs from MSTP instances.	
Syntax Description	mst-instance	MST instance. Range is 1-64.
	vlan <vlan-range>	A single VLAN or a a range of VLANs. The format is <vlan> or <from-vlan>-<to-vlan>.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mst 1 vlan 10-20`	
Related Commands		
Note	This command is available when global STP mode is set to MST	

spanning-tree mst revision

spanning-tree mst revision <number>
no spanning-tree mst revision

Configures the MSTP revision number.
The no form of the command sets the parameter to its default value.

Syntax Description	number	The MST revision number. Range is 0-65535.
Default	0	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mst revision 1`	
Related Commands		
Note	• The revision number is one of three parameters, along with the MST name and VLAN-to-instance map, that identify the switch's MST region • This command is available when global STP mode is set to MST	

spanning-tree mst name

spanning-tree mst name <name>
no spanning-tree mst name

Configures the MSTP name.
The no form of the command sets the parameter to its default value.

Syntax Description	name	MST name: Up to 32 characters.
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Role	admin	
Example	`switch (config)# spanning-tree mst name my-mst`	
Related Commands		
Note	• The name is one of three parameters, along with the MST revision number and VLAN-to-instance map, that identifies the switch's MST region • This command is available when global STP mode is set to MST	

spanning-tree mst root

spanning-tree mst <mst-instance> root <role>
no spanning-tree mst <mst-instance> root

Changes the bridge priority for the specified MST instance to the following values:
- Primary – 8192
- Secondary – 16384

The no form of the command sets the parameter to its default value.

Syntax Description	mst-instance	MSTP instance. Possible range is 1-64.
	role	Values: “primary” or “secondary”.
Default	primary	
Configuration Mode	config	
History	3.3.4150	
	3.7.1000	Updated example
Role	admin	
Example	`switch (config)# spanning-tree mst 1 root primary`	
Related Commands		
Note	• The root command is a way to automate a system configuration while ‘playing’ with the priority field. The priority field granularity may be too explicit for some users in case you wish to have 2 levels of priority (primary and secondary). So by default all the switches get the same priority and while using the root option you can get the role of master and backup by setting the priority field to a predefined value. • This command is available when global STP mode is set to MST.	

spanning-tree mst port-priority

spanning-tree mst {mst-instance} port-priority <priority>
no spanning-tree mode

Changes the spanning tree mode.
The no form of the command sets the parameter to its default value.

Syntax Description	mst-instance	MST instance. Range is 1-64.
	priority	MST instance port priority. Valid values are: 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224 and 240.
Default	rst	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.3.4150	
Role	admin	
Example	`switch (config  1/1)# spanning-tree mst 1 port-priority 32768` `switch (config interface port-channel 1)# spanning-tree mst 1 port-priority 32768`	
Related Commands		
Note	This command is available when global STP mode is set to MST.	

spanning-tree mst cost

	spanning-tree mst {mst-instance} cost <cost-value> **no spanning-tree mode** Configures the cost per MSTP instance. The no form of the command sets the parameter to its default value.	
Syntax Description	mst-instance	MST instance. Range is 1-64.
	cost-value	MST instance port cost. Range is 0-200000000.
Default	20000 for 1Gb/s, 2000 for 10Gb/s, 500 for 40Gb/s, 357 for 56Gb/s, 200 for 100Gb/s	
Configuration Mode	config interface port-channel	
History	3.3.4150	
Role	admin	
Example	`switch (config  1/1)# spanning-tree mst 1 cost 4000` `switch (config interface port-channel 1)# spanning-tree mst 1 cost 4000`	
Related Commands		
Note	This command is available when global STP mode is set to MST.	

spanning-tree vlan forward-time

spanning-tree vlan <vid> forward-time <secs>
no spanning-tree vlan <vid> forward-time

Configures how fast an interface changes its spanning tree state from Blocking to Forwarding.
The no form of the command resets the parameter value to its default.

Syntax Description	secs	Parameter range: 4-30 seconds.
Default	15 seconds	
Configuration Mode	config	
History	3.4.1100	
Role	admin	
Example	`switch (config) # spanning-tree vlan 10 forward-time 15`	
Related Commands	show spanning-tree	
Note	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

spanning-tree vlan hello-time

spanning-tree vlan <vid> hello-time <secs>
no spanning-tree vlan <vid> hello-time

Configures how often the switch broadcasts its hello message to other switches when it is the root of the spanning tree.
The no form of the command resets the parameter value to its default.

Syntax Description	secs	Parameter range: 1-2 seconds.
Default	2 seconds	
Configuration Mode	config	
History	3.4.1100	
Role	admin	
Example	`switch (config) # spanning-tree vlan 10 hello-time 2`	
Related Commands	show spanning-tree	
Note	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

spanning-tree vlan max-age

spanning-tree vlan <vid> max-age <secs>
no spanning-tree vlan <vid> max-age

Sets the maximum age allowed for the Spanning Tree Protocol information learned from the network on any port before it is discarded.
The no form of the command resets the parameter value to its default.

Syntax Description	secs	Parameter range: 6-40 seconds.
Default	20 seconds	
Configuration Mode	config	
History	3.4.1100	
Role	admin	
Example	`switch (config) # spanning-tree vlan 10 max-age 20`	
Related Commands	show spanning-tree	
Note	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

spanning-tree vlan priority

spanning-tree vlan <vid> priority <priority>
no spanning-tree vlan <vid> priority

Configures RPVST instance port priority.
The no form of the command resets the parameter value to its default.

Syntax Description	priority	Possible values are: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440.
Default	32768	
Configuration Mode	config	
History	3.4.1100	
Role	admin	
Example	`switch (config) # spanning-tree vlan 10 priority 32768`	
Related Commands	show spanning-tree	
Note	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

show spanning-tree

	show spanning-tree Displays spanning tree information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.1100 Updated Example with R and G flags 3.6.6000 Updated Example 3.6.6102 Added note on MLAG spanning-tree cost
Role	admin

Example	`switch (config) # show spanning-tree`
	
	`Switch                    : ethernet-default`
	`Spanning tree protocol rst : enabled`
	`Spanning tree force version: 2`
	
	`Root ID:`
	`  Priority: 32768`
	`  Address : 7c:fe:90:ff:2c:40`
	
	`  This bridge is the root`
	
	`  Hello Time (sec)   : 2`
	`  Max Age (sec)      : 20`
	`  Forward Delay (sec): 15`
	
	`Bridge ID:`
	`  Priority           : 32768`
	`  Address            : 7c:fe:90:ff:2c:40`
	`  Hello Time (sec)   : 2`
	`  Max Age (sec)      : 20`
	`  Forward Delay (sec): 15`
	
	`L: Loop Inconsistent`
	`R: Root Inconsistent`
	`G: BPDU Guard Inconsistent`
	
	`----------------------------------------------------------------------`
	`Interface        Role        Sts             Cost     Prio   Type`
	`----------------------------------------------------------------------`
	`Eth1/7           Designated  Discarding      200      128    normal`
	`Eth1/8           Disabled    Discarding(G)   200      128    edge`

Related Commands	clear spanning-tree counters spanning-tree
Note	MLAG spanning-tree cost is always equal to the cost of there being 2 member ports in the MLAG (even if one of the member ports fails or a new port is added).

show spanning-tree detail

	show spanning-tree detail Displays detailed spanning-tree configuration and statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.4110 Updated example output 3.6.5000 Updated example output
Role	admin

Example	`switch (config) # show spanning-tree detail`
	
	`Switch ethernet-default`
	`Spanning tree protocol is enabled`
	`Bridge is executing the rst compatible Spanning Tree Protocol`
	
	`Bridge Identifier priority: 32768`
	`Bridge Identifier address: e4:1d:2d:3d:5e:c0`
	`Configured hello time: 2, max age 20, forward delay 15`
	`Current root: priority 32768, address e4:1d:2d:3d:5e:c0`
	`Number of topology changes: 1, last change occurred 00:00:02 ago`
	`Last TCN received from: N/A`
	`Timers: hold 6 hello 2, max age 20, forward delay 15`
	`Default port type: normal`
	`Default bpdu filter: disabled`
	`Default bpdu guard: disabled`

Related Commands	clear spanning-tree counters spanning-tree
Note	

show spanning-tree interface

show spanning-tree interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel>

Display running state for specific interfaces.

Syntax Description	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
Role	admin	

Example	`switch (config) # show spanning-tree  1/2`
	`Eth1/2 is Disabled  Discarding`
	`  Port path cost 500, Port priority 128, Port Identifier 128.5`
	`  Designated root has priority 0, address unknown`
	`  Designated bridge has priority 0, address unknown`
	`  Designated port id 0.0, designated path cost 0`
	`  Number of transitions to forwarding state: 0`
	`  Port type: normal`
	`  PortFast is: off`
	`  Bpdu filter: disabled`
	`  Bpdu guard: disabled`
	`  Loop guard: disabled`
	`  Root guard: disabled`
	`  Link type: point-to-point`
	`  BPDU: sent: 0   received: 0`
	`switch (config) #`	

Related Commands	clear spanning-tree counters spanning-tree	
Note		

show spanning-tree mst

show spanning-tree mst [details | <instance> interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel>}]

Displays basic multi-spanning-tree information.

Syntax Description	details	Displays detailed multi-spanning-tree configuration and statistics
	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.6000	Updated Example
Role	admin	

Example	`switch (config) # show spanning-tree mst`
	
	`MST0:`
	`  vlans mapped: 1-1023,1025-2047,2049-3071,3073-4094`
	
	`L: Loop Inconsistent`
	`R: Root Inconsistent`
	`G: BPDU Guard Inconsistent`
	
	`---------------------------------------------------------------------`
	`Interface  Role  Sts  Cost  Prio  Type`
	`---------------------------------------------------------------------`
	`Eth1/7  Designated  Discarding  200  128.7  normal`
	`Eth1/8  Disabled  Discarding(G)  200  128.8  edge`	

Related Commands	clear spanning-tree counters spanning-tree	
Note		

show spanning-tree root

show spanning-tree root

Displays root multi-spanning-tree information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4150
Role	admin

Example	`switch (config) # show spanning-tree root`
	`Instance Priority MAC addr Root Cost Hello Time Max Age FWD Dly Root Port`
	`------- ------ -------- --------- -------- -------- ------- ---------`
	`MST0 32768 00:02:c9:71:ed:40 500 2 20 15 Eth1/20`
	`MST1 32768 00:02:c9:71:f0:c0 0 2 20 15 -`
	`MST2 0 00:02:c9:71:f0:c0 0 2 20 15 -`
	`MST3 32768 00:02:c9:71:f0:c0 0 2 20 15 -`
	`switch (config) #`

Related Commands	clear spanning-tree counters spanning-tree
Note	

show spanning-tree vlan

show spanning-tree vlan <vid> [detail | interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel>}]

Displays spanning-tree protocol information.

Syntax Description	vid	VLAN ID. Range is also supported Format: <vid1>[-<vid2>]
	detail	Displays detailed RPVST configuration and statistics
	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.5000	Updated example output
Role	admin	

Example	`switch (config) # show spanning-tree detail`
	
	`Switch ethernet-default`
	`Spanning tree protocol is enabled`
	`Bridge is executing the rpvst compatible Spanning Tree Protocol`
	
	`Vlan 1:`
	`  Bridge Identifier priority: 32769`
	`  Bridge Identifier address: e4:1d:2d:3d:5e:c0`
	`  Configured hello time: 2, max age 20, forward delay 15`
	`  Current root: priority 32769, address e4:1d:2d:3d:5e:c0`
	`  Number of topology changes: 0, last change occurred 00:00:00 ago`
	`  Last TCN received from: N/A`
	`  Timers: hold 6 hello 2, max age 20, forward delay 15`
	`  Default port type: normal`
	`  Default bpdu filter: disabled`
	`  Default bpdu guard: disabled`
	`switch (config) #`	

Related Commands	clear spanning-tree counters spanning-tree	
Note		

show spanning-tree vlan topo-change-history

show spanning-tree vlan <vid> topo-change-history

Displays spanning-tree topology change notification history per VLAN.

Syntax Description	vid	VLAN ID. Range is also supported. Format: <vid1>[-<vid2>]
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Role	admin	

Example	`switch (config) # show spanning-tree vlan 50 topo-change-history`
	
	`Vlan 50`
	
	`------------------------------------`
	`Interface  Date  Time`
	`------------------------------------`
	`Eth1/49  07/18/17  04:39:58`
	`Eth1/49  07/18/17  04:39:55`
	`Eth1/49  07/18/17  04:38:11`
	`Eth1/49  07/18/17  04:38:09`	

Related Commands	spanning-tree	
Note		

show spanning-tree mst topo-change-history

Show spanning-tree mst <mst-instance> topo-change-history

Displays spanning-tree topology change notification history per instance.

Syntax Description	mst-instance	MST instance. Range is 1-64.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Role	admin	

Example	`switch (config) # show spanning-tree mst 5 topo-change-history`
	
	`Instance 5`
	
	`------------------------------------`
	`Interface Date Time`
	`------------------------------------`
	`Eth1/49 07/18/17 04:43:51`
	`Eth1/49 07/18/17 04:43:33`	

Related Commands	spanning-tree	
Note		

show spanning-tree topo-change-history

show spanning-tree topo-change-history

Displays spanning-tree topology change notification history.

Syntax Description	mst-instance	MST instance. Range is 1-64.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Role	admin	
Example	`switch (config) # show spanning-tree topo-change-history`	

```
-----------------------------------
Interface    Date          Time
-----------------------------------
Eth1/49      07/27/17      09:39:38
Eth1/35      07/27/17      09:35:42
Eth1/35      07/27/17      09:35:40
Eth1/35      07/27/17      09:35:08
Eth1/35      07/27/17      09:35:06
Eth1/35      07/27/17      09:32:05
Eth1/35      07/27/17      09:32:03
Eth1/35      07/27/17      09:31:42
Eth1/35      07/27/17      09:31:40
```

Related Commands	spanning-tree
Note	

5.7 MAC Address Table

5.7.1 Configuring Unicast Static MAC Address

You can configure static MAC addresses for unicast traffic. This feature improves security and reduces unknown unicast flooding.

➢ ***To configure Unicast Static MAC address:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Run the command "`mac-address-table static unicast <destination mac address> vlan <vlan identifier(1-4094)> <slot>/ <port>`".

```
switch (config) # mac-address-table static 00:11:22:33:44:55 vlan 1  1/1
```

5.7.2 MAC Learning Considerations

MAC learning may be disabled using the command `mac-learning disable` which is beneficial in the following situations:

- To prevent denial-of-service attacks
- To manage the available MAC address table space by controlling which interfaces can learn MAC addresses
- To duplicate to a dedicated server (port7) all the packets that one host (host1; port1) sends to another (host2; port2), like in port mirroring. To accomplish this, MAC learning is disabled on port2. In this case the FDB does not obtain the MAC address of host2. Also, to prevent broadcast to every port, it is possible to configure a VLAN (VLAN 80) which ports 1, 2 and 7 are member of.

Figure 19: MAC Learning Disable Example Case

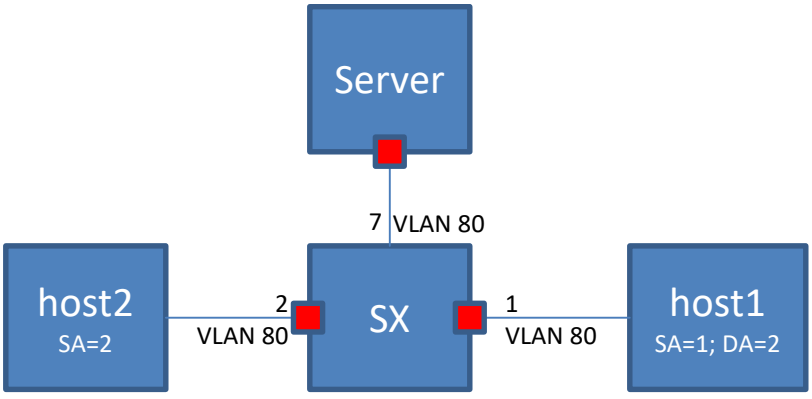

5.7.3 Commands

mac-address-table aging-time

mac-address-table aging-time <age>
no mac-address-table aging-time

Sets the maximum age of a dynamically learnt entry in the MAC address table.
The no form of the command resets the aging time of the MAC address table to its default.

Syntax Description	age	10-1000000 seconds.
Default	300	
Configuration Mode	config	
History	3.1.0600	
Role	admin	

Example	`switch (config) # mac-address-table aging-time 50`
	`switch (config) # show mac-address-table aging-time`
	
	`Mac Address Aging Time: 50`
	
	`switch (config) #`	

Related Commands	show mac-address-table show mac-address-table aging time	
Note		

mac-address-table static

mac-address-table static <mac address> vlan <vlan> interface <if-type> <if-number>
no mac-address-table static <mac address> vlan <vlan> interface <if-type> <if-number>

Configures a static MAC address in the forwarding database.
The no form of the command deletes a configured static MAC address from the forwarding database.

Syntax Description	mac address	Destination MAC address.
	vlan	VLAN ID or VLAN range.
	if-type	Ethernet or port-channel interface type.
	if-number	The interface number (i.e. 1/1, 3).
Default	No static MAC addresses available in default.	
Configuration Mode	config	
History	3.1.0600	
Role	admin	

Example	`switch (config) # mac-address-table static aa:aa:aa:aa:aa:aa vlan 1 interface ethernet 1/7`
	`switch (config) # show mac-address-table`
	
	`Switch ethernet-default`
	
	`Vlan  Mac Address  Type  Interface`
	`----  -----------  ----  ------------`
	`1  aa:aa:aa:aa:aa:aa  static  Eth1/7`
	`Number of unicast:  1`
	`Number of multicast:  0`
	`switch (config) #`	

Related Commands	show mac-address-table mac-address-table aging time	
Note	The no form of the command will not clear a dynamic MAC address. Dynamic MAC addresses are cleared using the “clear mac-address-table dynamic” command.	

mac-learning disable

	mac-learning disable **no mac-learning disable** Disables MAC-address learning. The no form of the command enables MAC-address learning.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface ethernet config interface port-channel
History	3.1.0600
Role	admin
Example	`switch (config  1/1) # interface ethernet mac-learning disable`
Related Commands	
Note	• When adding a port to a LAG, the port needs to be aligned with the LAG's configuration • When removing a port from a LAG, the port remains in whichever configuration the LAG is in • Disabling MAC learning is not supported on a local analyzer port. • Disabling MAC learning is not supported on an IPL LAG.

clear mac-address-table dynamic

clear mac-address-table dynamic

Clear the dynamic entries in the MAC address table.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0600
Role	admin
Example	`switch (config) # clear mac-address-table dynamic` `switch (config) #`
Related Commands	mac-address-table aging-time mac-address-table static show mac-address-table
Note	This command does not clear the MAC addresses learned on the mgmt0 port. Static entries are deleted using the “no mac-address-table static” command.

show mac-address-table

show mac-address-table [address <mac-address> | <if-number> | vlan [<vlan> | range <range>] | unicast | multicast]

Displays the static and dynamic unicast and multicast MAC addresses for the switch. Various of filter options available.

Syntax Description	mac-address	Filter the table to a specific MAC address.
	if-number	Filter the table to a specific interface.
	vlan	Filter the table to a specific VLAN number (1-4094).
	range	Filter the table to a range of VLANs.
	unicast	Filter the table to a unicast addresses only.
	multicast	Filter the table to a multicast addresses only.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0600	
	3.3.4500	Updated Example
Role	admin	

Example	`switch (config) # show mac-address-table`
	
	`Switch ethernet-default`
	
	`Vlan Mac Address Type Interface`
	`---- ----------- ---- ------------`
	`1 00:00:00:00:00:01 Static Po5`
	`1 00:00:3D:5C:FE:16 Dynamic Eth1/1`
	`1 00:00:3D:5D:FE:1B Dynamic Eth1/2`
	`Number of unicast: 2`
	`Number of multicast: 0`
	`switch (config) #`	

Related Commands	mac-address-table static clear mac-address-table	
Note		

show mac-address-table aging-time

	show mac-address-table aging-time Displays the MAC address table aging time.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0600
Role	admin

Example	`switch (config) # mac-address-table aging-time 300`
	`switch (config) # show mac-address-table aging-time`
	
	`Mac Address Aging Time: 300`
	
	`switch (config) #`

Related Commands	mac-address-table aging-time mac-address-table static clear mac-address-table
Note	MAC addresses learned on the mgmt0 is not shown by this command.

show mac-address-table interface

show mac-address-table interface [port-channel | mlag-port-channel <if>]

Displays the MAC address table of a port channel or an MLAG port channel.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4006
Role	admin

Example	`switch (config) # show mac-address-table`
	`----------`
	`Vlan Mac Address Type Port`
	`----------`
	`1 E4:1D:2D:37:11:22 Static Eth1/1`
	`1 E4:1D:2D:37:3E:11 Static Po5`
	
	`Number of unicast: 2`
	`Number of multicast: 0`
	
	`switch (config) # show mac-address-table interface port-channel 5`
	`----------`
	`Vlan Mac Address Type Port`
	`----------`
	`1 E4:1D:2D:37:3E:11 Static Po5`
	
	`Number of unicast: 1`
	`Number of multicast: 0`

Related Commands	mac-address-table static clear mac-address-table
Note	

show mac-address-table summary

	show mac-address-table summary Displays total number of unicast/multicast MAC address entries.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Role	admin

Example	`switch (config) # show mac-address-table summary`
	`Number of unicast:    4`
	`Number of multicast:  0`

Related Commands	mac-address-table static clear mac-address-table
Note	

5.8 MLAG

A link aggregation group (LAG) is used for extending the bandwidth from a single link to multiple links and provide redundancy in case of link failure. Extending the implementation of the LAG to more than a single device provides yet another level of redundancy that extends from the link level to the node level. This extrapolation of the LAG from single to multiple switches is referred to as multi-chassis link aggregation (MLAG). MLAG is supported on Ethernet blades' internal as well as external ports.

Figure 20: Basic MLAG Setup

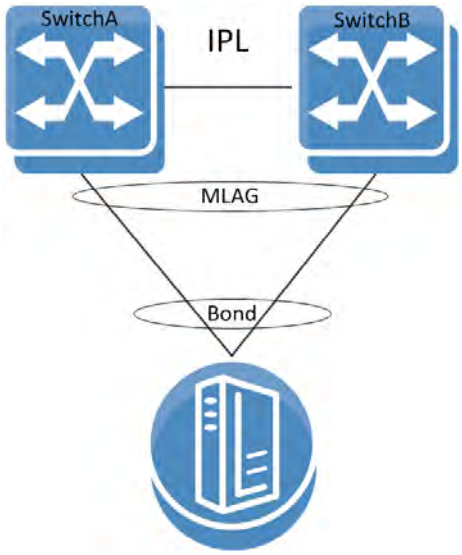

Each switch configuration is independent and it is user responsibility to make sure to configure both switches similarly pertaining MLAG (e.g. MLAG port-channel VLAN membership, static MAC, ACL, etc).

A peered device (host or switch) connecting to switches running an MLAG runs a standard LAG and is unaware of the fact that the LAG connects to two separate switches.

The MLAG switches share an inter-peer link (IPL) between them for carrying control messages in a steady state or data packages in failure scenarios. Thus, the bandwidth of the IPL should be defined accordingly. The IPL itself can be a LAG and may be constructed of links of any supported speed. In such a case, PFC must be configured on this IPL. Figure 21, "Basic MLAG Topology," on page 857 illustrates this. The IPL serves the following purposes:

- MLAG protocol control – keepalive messages, MAC sync, MLAG port sync, etc.
- MLAG port failure – serves redundancy in case of a fallen link on one of the MLAG switches
- Layer-3 failure – serves redundancy in case of a failed connection between the MLAG switches and the rest of the L3 network should there be one

The IPL VLAN interface must be used only for MLAG protocol and must not be used by any other interfaces (e.g. LAG, Ethernet).

The MLAG protocol is made up of the following components to be expanded later:

- Keepalive
- Unicast and multicast sync
- MLAG port sync

When positioned at the top of rack (ToR) and connecting with a Layer-3 uplink, the MLAG pair acts as the L3 border for the hosts connected to it. To allow default gateway redundancy, both MLAG switches should be addressed by the host via the same default gateway address.

MLAG uses an IP address (VIP) that points to all MLAG member nodes.

When running MLAG as L2/L3 border point, an MAGP VIP must be deployed as the default GW for MLAG port-channels (MPOs).

When MLAG is connected through a Layer-2 based uplink, there is no need to apply default gateway redundancy towards hosts since this function is implemented on the L2/L3 border points of the network.

For more information, refer to Section 6.9, “MAGP,” on page 1572.

The two peer switches need to carry the exact same configuration of the MLAG attributes for guaranteeing proper functionality of the MLAG.

Ensuring that both switches are configured identically is the responsibility of the user and is not monitored by the Onyx software.

MLAG is currently supported for 2 switches only.

The VIP address must be on the same management IP subnet.

All nodes in an MLAG must be of the same CPU type (e.g. x86), switch type (e.g. Spectrum), and must all have the same OS version installed.

When working with MLAG, the maximum number of MAC addresses is limited to 88K. Without it, there is no limitation.

When transitioning from standalone into a group or vice versa, a few seconds are required for the node state to stabilize. During that time, group feature commands (e.g. MLAG commands) should not be executed. To run group features, wait for the CLI prompt to turn into [standalone:master], [<group>:master] or [<group>:standby] instead of [standalone:*unknown*] or [<group>:*unknown*].

Each MLAG VIP group must be configured with a different unicast IP address. If not, MLAG behavior is inconsistent.

In a scenario where there is no IP communication between the MGMT ports of the MLAG switches (for example when one MGMT port is disconnected), the following CLI prompt is displayed: `<hostname>[<mlag cluster name>:unknown]#`. This does not reflect the MLAG state, but only the state of the cluster.

5.8.1 MLAG Keepalive and Failover

Master election in MLAG is based on the IPs of the nodes taking part of the MLAG. The master elected is that which has the highest IPL VLAN interface local IP address.

MLAG master/slave roles take effect in fault scenarios such as split-brain, peer faults, and during software upgrades.

The MLAG pair of switches periodically exchanges a keepalive message on a user configurable interval. If the keepalive message fails to arrive for three consecutive intervals the switches break into two standalone switches. In such a case, the remaining active switch begins to act as a standalone switch and assumes that its previously peering MLAG switch has failed.

To avoid a scenario where failure on the IPL causes both MLAG peers to assume that their peer has failed, a safety mechanism is maintained based on UDP packets running via the management plane which alerts both MLAG switches that its peer is alive. In such case where keepalive packets are not received the slave shuts down its MLAG interfaces and the master becomes a standalone switch in order to avoid misalignment in MLAG configuration.

5.8.2 Unicast and Multicast Sync

Unicast and multicast sync is a mechanism which syncs the unicast and multicast FDBs of the MLAG peers. It prevents unicast asymmetric traffic from loading the network with flood traffic and multicast traffic from being processed.

5.8.3 MLAG Port Sync

Under normal circumstances, traffic from the IPL cannot pass through the MLAG ports (the IPL is isolated from the MLAG ports). If one of the MLAG links break, the other MLAG switch opens that isolation and allows traffic from its peer through the IPL to flow via the MLAG port which accesses the destination of the fallen link.

5.8.4 MLAG Virtual System-MAC

A pair of MLAG switches uses a single virtual system MAC for L2 protocols (such as LACP) operating on the MLAG ports. This virtual system MAC is served also as the STP bridge ID.

The virtual system MAC is automatically computed based on the MLAG VIP name, but can be manually set using the command "system-mac".

MLAG relies on systems to have the same virtual system MAC. Therefore, if a system MAC mismatch is detected, the slave shuts down its interfaces.

5.8.5 Upgrading MLAG Pair

Switches in the same MLAG group must have the same Onyx version.

When peers identify having different versions, they enter an upgrading state in which the slave peer waits for a specific period of time (according to the command "upgrade-timeout" on page 873) before closing its ports.

It is advised to plan MLAG upgrade in advance and perform it in a timely manner. Please avoid performing topology changes during the upgrade period.

For more information on MLAG upgrade, please see Section 4.5.3, "Upgrading Onyx HA Groups," on page 278.

When two tiers of MLAG pairs are used, each pair should be upgraded sequentially and not in parallel to prevent traffic loops.

5.8.6 Interoperability with MLAG

5.8.6.1 MLAG Interoperability with L2 Protocols

MLAG inter-operates with all STP modes (RSTP, MSTP and PVRST). MLAG can be configured in a spanning tree network where the two MLAG switches function as one STP entity.

In general all static configuration must be configured identically on both peers.

Table 52 - L2 Protocol Interoperability with MLAG

Protocol	Description
Static MAC addresses	Static MAC address are not synced between MLAG peers
LACP	MPO supports all LACP modes (passive/active), but it is not a must. If used, their configuration must be identical on each peer.

Table 52 - L2 Protocol Interoperability with MLAG

Protocol	Description
VLAN	VLAN membership of an MPO must be configured identically on both peers. This includes PVID, switchport mode, and tagged/untagged VLAN. VLAN static configuration such as snooping MRouter must be configured identically on both peers as well.
Spanning-tree protocol	MPO spanning-tree configuration must be identical in both switches, and other local ports' spanning-tree configuration must be done when those ports are down.
IGMP snooping	IGMP snooping must be activated globally on both peers. IGMP snooping attributes on the MPO must have identical configuration.
Port mirroring	Supported
PIM	Not supported
sFlow	Supported
LLDP	All attributes of a the MPO must be configured identically on both peers.
Isolation-groups	Not supported with MLAG
OpenFlow	Not supported over MLAG IPL
PTP	Not supported over MLAG IPL (not supported over LAG in general)
NVE	Not supported
Dot1x	Not supported

5.8.6.2 MLAG Interoperability with L3 Protocols

Onyx cannot route between the two MLAG switches. That is, the source and the client cannot be connected to MPOs at the same time (only one at a time).

For cases when we need to re-direct the traffic, another physical link is needed which is not part of the IPL (preferably a router port) to connect the two switches.

Dynamic routing protocols (e.g. OSPF, BGP) are not supported over MPOs. If they are necessary, router ports must be used instead of MPOs.

5.8.7 Configuring MLAG

This section provides a basic example of how to configure two switches and a server in an MLAG setup.

Figure 21: Basic MLAG Topology

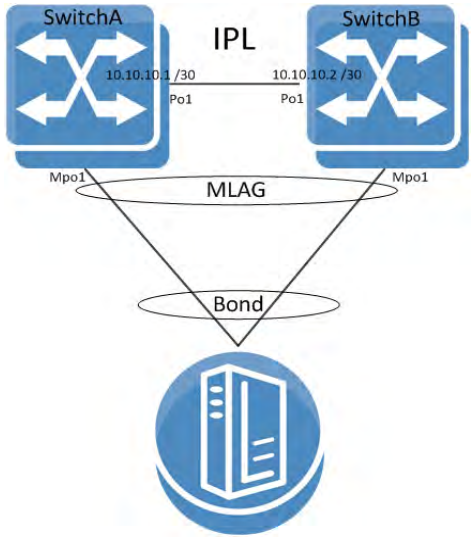

For more advanced configuration options, please refer to the following Mellanox Community post: https://community.mellanox.com/docs/DOC-2262.

➢ ***To configure L2 MLAG:***

Prerequisites:

Step 1. Enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. (Recommended) Enable LACP in the switch. Run:

```
switch (config)# lacp
```

Step 3. Enable QoS on the switch to avoid congestion on the IPL port. Run:

```
switch (config)# dcb priority-flow-control enable force
```

Step 4. Enable the MLAG protocol commands. Run:

```
switch (config)# protocol mlag
```

Configuring the IPL:

Step 1. Create a VLAN for the inter-peer link (IPL) to run on. Run:

```
switch (config)# vlan 4000
switch (config vlan 4000)#
```

Step 2. Create a LAG. Run:

```
switch (config)# interface port-channel 1
switch (config interface port-channel 1)#
```

Step 3. Map a physical port to the LAG in active mode (LACP). Run:

```
switch (config)# interface ethernet 1/1 channel-group 1 mode active
```

Step 4. Set this LAG as an IPL. Run:

```
switch (config interface port-channel 1)# ipl 1
```

Step 5. Enable QoS on this specific interface. Run:

```
switch (config interface port-channel 1)# dcb priority-flow-control mode on force
```

Step 6. Create a VLAN interface. Run:

```
switch (config)# interface vlan 4000
switch (config interface vlan 4000)#
```

Step 7. Set an IP address and netmask for the VLAN interface.
Configure IP address for the IPL link on both switches:
Note: The IPL IP address should not be part of the management network, it could be any IP address and subnet that is not in use in the network. This address is not advertised outside the switch.

On SwitchA, run:

```
switch (config interface vlan 4000)# ip address 1.1.1.1 /30
```

On SwitchB, run:

```
switch (config interface vlan 4000)# ip address 1.1.1.2 /30
```

Step 8. Map the VLAN interface to be used on the IPL and set the peer IP address (the IP address of the IPL port on the second switch) of the IPL peer port. IPL peer ports must be configured on the same netmask.

On SwitchA, run:

```
switch (config interface vlan 4000)# ipl 1 peer-address 1.1.1.2
```

On SwitchB, run:

```
switch (config interface vlan 4000)# ipl 1 peer-address 1.1.1.1
```

Step 9. (Optional) Configure a virtual IP (VIP) for the MLAG. MLAG VIP is important for retrieving peer information.

Note: If you have a mgmt0 interface, the IP address should be within the subnet of the management interface. Do not use mgmt1. The management network is used for keepalive messages between the switches. The MLAG domain must be unique name for each MLAG domain. In case you have more than one pair of MLAG switches on the same network, each domain (consist of two switches) should be configured with different name.

On SwitchA, run:

```
switch (config)# mlag-vip my-vip ip 10.234.23.254 /24
```

On SwitchB, run:

```
switch (config)# mlag-vip my-vip
```

Step 10. (Optional) Configure a virtual system MAC for the MLAG. Run:

```
switch (config)# mlag system-mac 00:00:5E:00:01:5D
```

Creating an MLAG interface:

Step 1. Create an MLAG interface for the host. Run:

```
switch (config)# interface mlag-port-channel 1
switch (config interface mlag-port-channel 1)#
```

Step 2. Bind an Ethernet port to the MLAG group. Run:

```
switch (config  1/2)# mlag-channel-group 1 mode on
```

Step 3. Create and enable the MLAG interface. Run:

```
switch (config interface mlag-port-channel 1)# no shutdown
```

Enabling MLAG:

Step 1. Enable MLAG. Run:

```
switch [my-vip: master] (config mlag)# no shutdown
```

When running MLAG as L2/L3 border point, MAGP VIP must be deployed as the default GW for MPOs. For more information, refer to Section 6.9, “MAGP,” on page 1572.

➢ ***To verify MLAG configuration:***

Step 1. Examine MLAG configuration and status. Run:

```
SX2 [mellanox: master] (config)# show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 1 sec
Keepalive-interval: 30 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5E:00:01:5D

MLAG Ports Configuration Summary:
Configured: 1
 Disabled:   0
 Enabled:    1

MLAG Ports Status Summary:
Inactive:       0
 Active-partial: 0
 Active-full:    1
```

```
MLAG IPLs Summary:
ID   Group         Vlan       Operational    Local         Peer        Up Time     Toggle Counter
    Port-Channel  Interface  State        IP address    IP address
-------------------------------------------------------------------------------------------
1    Po1           1          Up          10.10.10.1     10.10.10.2   0 days      00:00:09 5
Peers state Summary:
System-id          State   Hostname
----------------------------------
F4:52:14:2D:9B:88  Up      <SX2>
F4:52:14:2D:9B:08  Up       SX1
switch [mellanox: master] (config)#
```

Step 2. Examine the MLAG summary table. Run:

```
switch [my-vip: master] (config)# show interfaces mlag-port-channel summary

MLAG Port-Channel Flags: D-Down, U-Up, P-Partial UP, S-suspended by MLAG

Port Flags:
  D: Down
  P: Up in port-channel (members)
  S: Suspend in port-channel (members)
  I: Individual

MLAG Port-Channel Summary:
  ---------------------------------------------------------------------------
  Group              Type     Local                    Peer
  Port-Channel                Ports                    Ports
  (D/U/P/S)                   (D/P/S/I)                (D/P/S/I)
  ---------------------------------------------------------------------------
  1 Mpo2(U)          Static   Eth1/2(P)                Eth1/2(P)

switch (config)#
```

Step 3. Examine the MLAG statistics. Run:

```
switch [my-vip: master] (config)# show mlag statistics
IPL 1:
  Rx Heartbeat           : 516
  Tx Heartbeat           : 516
  Rx IGMP tunnel         : 0
  Tx IGMP tunnel         : 0
  RX XSTP tunnel         : 0
  TX XSTP tunnel         : 0
  RX mlag-notification   : 0
  TX mlag-notification   : 0
  Rx port-notification   : 0
  Tx port-notification   : 0
  Rx FDB sync            : 0
  Tx FDB sync            : 0
  RX LACP manager        : 1
  TX LACP manager        : 0
```

Enabling L3 Forwarding with User VRF

If you want to use a VRF for IP routing and forwarding on an MLAG topology, it is recommended to configure an additional VLAN interface with the same user VRF context as the non-MLAG L3 interface that has to route through the same physical ports as the IPL. This would allow forwarding L3 traffic through this VLAN interface on the same ports as the IPL.

5.8.8 Commands

protocol mlag

protocol mlag
no protocol mlag

Enables MLAG functionality and unhides the MLAG commands.
The no form of the command hides the MLAG commands and deletes its database.

Syntax Description	
Default	no protocol mlag
Configuration Mode	config
History	3.3.4500
Role	admin
Example	`switch (config) # protocol mlag` `switch (config) #`
Related Commands	
Note	• Running the no form of this command hides MLAG commands. • MLAG may be enabled without IP routing, but without IP routing an IPL vLAN interface cannot be configured and thus MLAG does not function. • MLAG may be enabled without IGMP snooping, but if IGMP snooping is disabled, multicast FDBs do not sync.

mlag

	mlag Enters MLAG configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4500
Role	admin
Example	`switch (config) # mlag` `switch (config mlag) #`
Related Commands	protocol mlag
Note	

shutdown

	shutdown **no shutdown** Disables MLAG. The no form of the command enables MLAG.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config mlag
History	3.3.4500
Role	admin
Example	`switch (config mlag) # no shutdown` `switch (config mlag) #`
Related Commands	protocol mlag
Note	This parameter must be similar in all MLAG peers.

interface mlag-port-channel

interface mlag-port-channel <if-number>
no interface mlag-port-channel <if-number>

Creates an MLAG interface.
The no form of the command deletes the MLAG interface.

Syntax Description	if-number	Integer. Interface number range: 1-1000.
Default	N/A	
Configuration Mode	config	
History	3.3.4500	
Role	admin	
Example	`switch (config) # interface mlag-port-channel 1` `switch (config interface mlag-port-channel 1) #`	
Related Commands	protocol mlag	
Note	• The maximum number of interfaces is 64. • The default Admin state is disabled. • Range configuration is possible on this interface. • This interface number must be the same in all the MLAG switches.	

ipl

ipl <ipl-id>
no ipl <ipl-id>

Sets this LAG as an IPL port.
The no form of the command resets this LAG as regular LAG.

Syntax Description	ipl-id	IPL ID. Only “1” IPL port is supported.
Default	no ipl	
Configuration Mode	config interface port-channel	
History	3.3.4500	
Role	admin	
Example	`switch (config interface port-channel 1)# ipl 1`	
Related Commands	protocol mlag	
Note	• If a LAG is set as IPL, only the commands “[no] shutdown”, “no ipl” and “no interface port-channel” become applicable. • A LAG interface set as IPL must have default LAG configuration, otherwise the set is rejected. Force option can be used.	

ipl peer-address

ipl <ipl-id> peer-address <IP-Address>
no ipl <ipl-id>

Maps a VLAN interface to be used for an IPL LAG and sets the peer IP address of the IPL peer port.
The no form of the command deletes a peer IPL LAG and unbinds this VLAN interface from the IPL function.

Syntax Description	ipl-id	IPL ID. Only "1" IPL port is supported.
	IP-Address	IPv4 address
Default	N/A	
Configuration Mode	config interface vlan	
History	3.3.4500	
Role	admin	
Example	`switch (config interface vlan 1)# ipl 1 peer-address 10.10.10.10` `switch (config interface vlan 1)#`	
Related Commands	protocol mlag	
Note	• The subnet mask is the same subnet mask of the VLAN interface. • This VLAN interface should be used for IPL only.	

keep-alive-interval

keep-alive-interval <value>
no keep-alive-interval

Configures the interval during which keep-alive messages are issued between the MLAG switches.
The no form of the command resets this parameter to its default value.

Syntax Description	value	Time in seconds. Range: 1-300.
Default	1 second	
Configuration Mode	config mlag	
History	3.3.4500	
Role	admin	
Example	`switch (config mlag) # keep-alive-interval 1` `switch (config mlag) #`	
Related Commands	protocol mlag	
Note	This parameter must be similar on all MLAG peers	

mlag-channel-group mode

mlag-channel-group <if-number> mode {on | active | passive}
no mlag-channel-group

Binds an Ethernet port to the MLAG port-channel (MPO).
The no form of the command deletes the binding.

Syntax Description	if-number	Integer. Interface number range: 1-1000.
	on	Binds to static MLAG
	active	Sets MLAG LAG in LACP active mode
	passive	Sets MLAG LAG in LACP passive mode
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.3.4500	
Role	admin	
Example	`switch (config  1/1)# mlag-channel-group 1 mode on` `switch (config  1/1)#`	
Related Commands	protocol mlag	
Note		

mlag-vip

mlag-vip <domain-name> ip [<ip-address> {<masklen> | netmask> [force]]
no mlag-vip

Sets the VIP domain and IP address for MLAG.
The no form of the command deletes the VIP domain and IP address.

Syntax Description	domain-name	MLAG group name
	<ip-address>	IP address
	<masklen>	Format example: /24. Note that a space is required between the IP address and the mask.
	<netmask>	Format example: 255.255.255.0. Note that a space is required between the IP address and the mask.
	force	Forces the IP address if another IP is already configured.
Default	N/A	
Configuration Mode	config	
History	3.3.4500	
Role	admin	
Example	`switch (config)# mlag-vip my-mlag-domain ip 10.10.10.254/24` `switch (config)#`	
Related Commands		
Note	• This IP address must be configured in one of the MLAG switches and must be in the box management subnet. • Other switches in the MLAG must join the same domain name.	

reload-delay

reload-delay <value>
no reload-delay

Specifies the amount of time that MLAG ports are disabled after system reboot. The no form of the command resets this parameter to its default value.

Syntax Description	value	Time in seconds. Range: 0-300.
Default	30 seconds	
Configuration Mode	config mlag	
History	3.3.4500	
Role	admin	
Example	`switch (config mlag) # reload-delay 30` `switch (config mlag) #`	
Related Commands		
Note	• This interval allows the switch to learn the IPL topology to identify the master and sync the MAC address before opening the MLAG ports. • This parameter must be similar in all MLAG peers.	

system-mac

system-mac <virtual-mac>
no system-mac <virtual-mac>

Configures virtual system MAC.
The no form of the command resets this value to its default value.

Syntax Description	virtual-mac MAC address
Default	Default is calculated according to the MLAG-VIP name, using the base MAC as VRRP MAC prefix (00:00:5E:00:01:xx) with the suffix hashed from the mlag-vip name 0...255.
Configuration Mode	config mlag
History	3.4.2008
Role	admin
Example	`switch (config mlag) # system-mac 00:00:5E:00:01:5D` `switch (config mlag) #`
Related Commands	
Note	This parameter must be configured the same in all MLAG peers.

upgrade-timeout

upgrade-timeout <time>
no upgrade-timeout

Configures the time period during which an MLAG slave keeps its ports active while in upgrading state.
The no form of the command resets the parameter value to its default.

Syntax Description	time	Time in minutes. Range: 0-120 minutes.
Default	60	
Configuration Mode	config mlag	
History	3.4.2008	
Role	admin	
Example	`switch (config mlag) # upgrade-timeout 60` `switch (config mlag) #`	
Related Commands		
Note	This parameter must be configured the same in all MLAG peers.	

show mlag

show mlag

Displays MLAG configuration and status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4500
	3.3.5006 — Updated example
	3.4.2008 — Updated example with system MAC and upgrade timeout
	3.6.5000 — Added note
	3.6.6102 — Updated example output
Role	admin
Example	

```
SX2 [mellanox: master] (config)# show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 1 sec
Keepalive-interval: 30 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5E:00:01:5D

MLAG Ports Configuration Summary:
Configured: 1
 Disabled:   0
 Enabled:    1

MLAG Ports Status Summary:
Inactive:       0
 Active-partial: 0
 Active-full:    1

MLAG IPLs Summary:
ID   Group        Vlan      Operational    Local         Peer        Up Time     Toggle Counter
    Port-Channel  Interface  State        IP address    IP address
-------------------------------------------------------------------------------------------
1    Po1          1          Up          10.10.10.1     10.10.10.2   0 days      00:00:09 5

MLAG Members Summary:
System-id          State   Hostname
-----------------------------------
F4:52:14:2D:9B:88  Up      <SX2>
F4:52:14:2D:9B:08  Up       SX1
SX2 [mellanox: master] (config)#
```

Related Commands	
Note	If run in the middle of an upgrade, the following message will appear in the output: `*Upgrading* <hostname> --> *Cluster upgrade in progress*`

show mlag-vip

show mlag-vip

Displays MLAG VIP configuration and status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4500 3.6.6102 Updated Example
Role	admin

Example	`switch (config)# show mlag-vip`
	`MLAG-VIP`
	` MLAG group name: Test`
	` MLAG VIP address: 10.10.10.3/24`
	` Active nodes: 2`
	
	`---------------------------------------------------------------`
	`Hostname VIP-State IP Address`
	`---------------------------------------------------------------`
	`SwitchA master 10.10.10.1`
	`SwitchB standby 10.10.10.2`

Related Commands	
Note	

show interfaces mlag-port-channel

show interfaces mlag-port-channel [<if-number>]

Displays the MLAG LAG configuration and status.

Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.1002	Added “error packets” counter to Tx
	3.6.5000	Added telemetry to output
	3.6.6000	Added “forwarding mode” to output
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config)# show interfaces mlag-port-channel 1

Mpo1:
  Admin state         : Disabled
  Operational state   : Down
  Description         : N\A
  Mac address         : N\A
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : 0 Gbps
  Auto-negotiation    : N/A
  Width reduction mode: Not supported
  Switchport mode     : access
  MAC learning mode   : Enabled
  Forwarding mode     : inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                  : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec
```

```
Rx:
   0                    packets
   0                    unicast packets
   0                    multicast packets
   0                    broadcast packets
   0                    bytes
   0                    discard packets
   0                    error packets
   0                    fcs errors
   0                    undersize packets
   0                    oversize packets
   0                    pause packets
   0                    unknown control opcode
   0                    symbol errors
   0                    discard packets by storm control

 Tx:
   0                    packets
   0                    unicast packets
   0                    multicast packets
   0                    broadcast packets
   0                    bytes
   0                    discard packets
   0                    error packets
   0                    hoq discard packets
```

Related Commands

Note

show interfaces mlag-port-channel counters

show interfaces mlag-port-channel <if-number> counters

Displays the extended counters for the interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.1002
Role	admin

Example	`switch (config)# show interfaces mlag-port-channel 3 counters`
	
	`Rx`
	`  12  packets`
	`  0  unicast packets`
	`  12  multicast packets`
	`  0  broadcast packets`
	`  2700  bytes`
	`  0  packets of 64 bytes`
	`  0  packets of 65-127 bytes`
	`  12  packets of 128-255 bytes`
	`  0  packets of 256-511 bytes`
	`  0  packets of 512-1023 bytes`
	`  0  packets of 1024-1518 bytes`
	`  0  packets Jumbo`
	`  0  error packets`
	`  0  discard packets`
	`  0  fcs errors`
	`  0  undersize packets`
	`  0  oversize packets`
	`  0  pause packets`
	`  0  unknown control opcode`
	`  0  symbol errors`
	
	`Tx`
	`  0  packets`
	`  0  unicast packets`
	`  0  multicast packets`
	`  0  broadcast packets`
	`  152100000000  bytes`
	`  100000000  error packets`
	`  0  discard packets`
	`  0  pause packets`
	`switch (config)#`

Related Commands	
Note	

show interfaces mlag-port-channel summary

show interfaces mlag-port-channel summary

Displays MLAG summary table.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4500 3.4.0000 Added notes and updated Example 3.4.1100 Updated Example 3.6.6000 Updated Example
Role	admin

Example	`switch [my-vip: standby] (config)# show interfaces mlag-port-channel summary`
	
	`MLAG Port-Channel Flags: D-Down, U-Up, P-Partial UP, S-suspended by MLAG`
	
	`Port Flags:`
	`D: Down`
	`P: Up in port-channel (members)`
	`S: Suspend in port-channel (members)`
	`I: Individual`
	
	`MLAG Port-Channel Summary:`
	`----------------------------------------------------------------------`
	`Group Type Local Peer`
	`Port-Channel Ports Ports`
	`(D/U/P/S) (D/P/S/I) (D/P/S/I)`
	`----------------------------------------------------------------------`
	`1 Mpo61(D) LACP Eth1/4(I) Eth1/3(S)`

Related Commands	
Note	• If a cluster is not available, the column “Peer Ports” shows “N/A”. If the cluster is available but is not configured on the peer, the “Peer Ports” column shows nothing. • If the system happens to be busy, peer ports may be unavailable and the following prompt may appear in the output: “System busy and partial information is presented – please try again later”. • The “I” flag indicates an interface which is part of a port-channel and in individual state • The “S” flag indicates an interface which is part of a port-channel and in suspended state

show mlag statistics

	show mlag statistics Displays the MLAG IPL counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4500 3.4.0000 Updated Example 3.6.6102 Updated Example
Role	admin

Example	`switch (config)# show mlag statistics`
	`IPL 1:`
	`  Rx Heartbeat  : 516`
	`  Tx Heartbeat  : 516`
	`  Rx IGMP tunnel  : 0`
	`  Tx IGMP tunnel  : 0`
	`  RX XSTP tunnel  : 0`
	`  TX XSTP tunnel  : 0`
	`  RX mlag-notification  : 0`
	`  TX mlag-notification  : 0`
	`  Rx port-notification  : 0`
	`  Tx port-notification  : 0`
	`  Rx FDB sync  : 0`
	`  Tx FDB sync  : 0`
	`  RX LACP manager  : 1`
	`  TX LACP manager  : 0`

Related Commands	
Note	

5.9 Link State Tracking

A group of links may contain upstream links and downstream links. When all upstream links in a group are down, Link State Tracking (LST) shuts all the downstream links down. In order to let the peer on the other side know that it needs to stop sending traffic on the downstream links. When the upstream link recovers, LST brings up the downstream links, letting the peers know that they may resume forwarding traffic on those links.

A link can be a member of several groups. A downstream interface is shut down if at least one of the groups requests a shutdown and is brought back up if all groups request it to be up.

In cases with only downstream links in a group (no upstream links), the downstream links will stay up.

5.9.1 Configuring Link State Tracking

The following is a basic example of how to configure link state tracking group and tracking VLAN.

Figure 22: Upstream interfaces

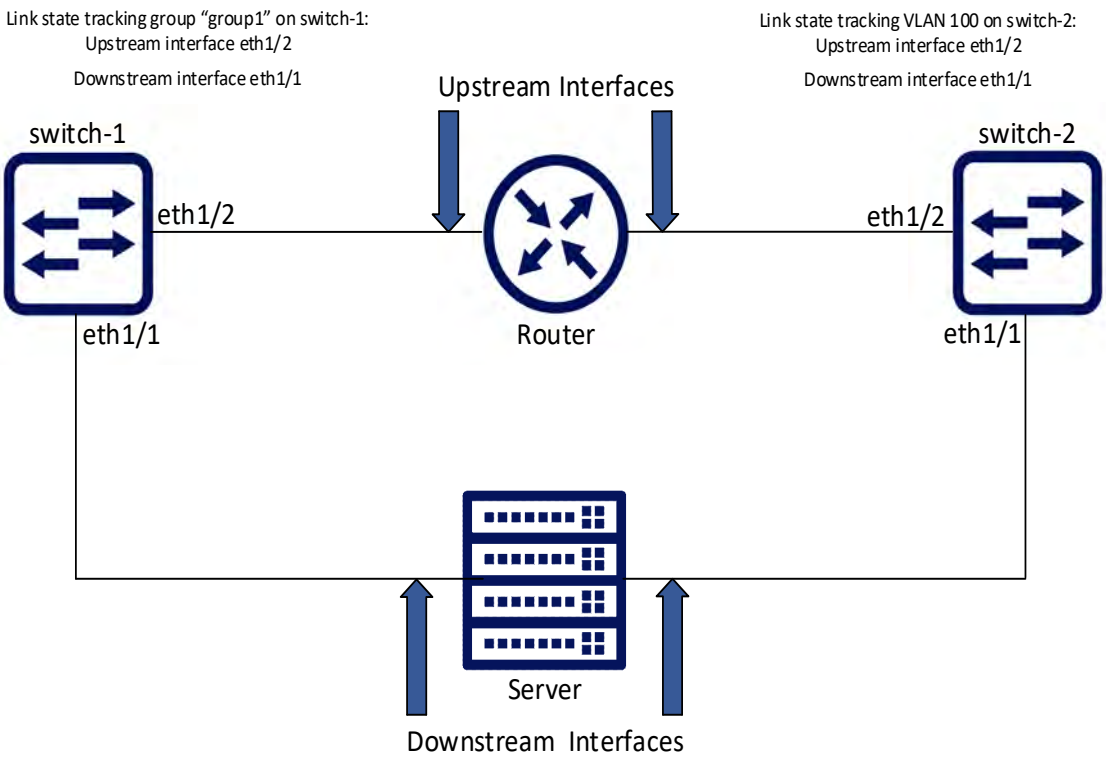

➢ ***To configure Link State Tracking group:***

Step 1. Create tracking group. Run:

```
switch-1 (config) # link state tracking group group1
```

Step 2. Configure link type on the interface. Run:

```
switch-1 (config) # interface ethernet 1/2 link type upstream
switch-1 (config) # interface ethernet 1/1 link type downstream
```

Step 3. Add interfaces into the group. Run:

```
switch-1 (config) # interface ethernet 1/1 link state tracking group group1
switch-1 (config) # interface ethernet 1/2 link state tracking group group1
```

➢ ***To configure Link State Tracking vlan:***

Step 1. Create vlan. Run:

```
switch-2 (config) # vlan 100
```

Step 2. Configure vlan members. Run:

```
switch-2 (config) # interface ethernet 1/1 switchport access vlan 100
switch-2 (config) # interface ethernet 1/2 switchport access vlan 100
```

Step 3. Configure link type on the interface. Run:

```
switch-2 (config) # interface ethernet 1/2 link type upstream
switch-2 (config) # interface ethernet 1/1 link type downstream
```

Step 4. Create link state tracking vlan. Run:

```
switch-2 (config) # link state tracking vlan 100
```

➢ ***To verify Link State Tracking configuration:***

```
switch-1 (config) # show link state tracking group group1

--------------------------------------------------------------------------------
Group              Port Type      Interface          Admin Status   Operational Status
--------------------------------------------------------------------------------
group1             Upstream       Eth1/2             Enabled        Up
group1             Downstream     Eth1/1             Enabled        Up


switch-2 (config) # show link state tracking vlan 100

--------------------------------------------------------------------------------
Group              Port Type      Interface          Admin Status   Operational Status
--------------------------------------------------------------------------------
Vlan 100           Upstream       Eth1/2             Enabled        Down
Vlan 100           Downstream     Eth1/1             Enabled        Down (by tracking)
```

5.9.2 Commands

link type

link type {downstream | upstream}
no link type

Configures an interface's link direction.
The no form of the command deletes the interface's link direction configuration.

Syntax Description	downstream	Configures interface as downstream
	upstream	Configures interface as upstream
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.7.1000	
Role	admin	
Example	`switch (config interface ethernet 1/1)# link type downstream`	
Related Commands	show link state tracking	
Note	• IPL, loopback, and VLAN interfaces are not supported. • An interface can be either upstream or downstream but not both.	

link state tracking group

link state tracking group <group-name>
no link state tracking group <group-name>

Creates a link state tracking group if one does not exist, and if applied to a specific interface, then it adds that interface to the group.
The no form of the command deletes a link state tracking group, and if applied to a specific interface, then it removes that interface from the group.

Syntax Description	group-name	group name
Default	N/A	
Configuration Mode	config config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.7.1000	Example updated
	3.7.11xx	Example updated
Role	admin	
Example	`switch (config)# link state tracking group group1` `switch (config interface ethernet 1/1)# link state tracking group group1`	
Related Commands		
Note	• The maximum number of tracking groups/vlans is 64 • Link state tracking group name should not contain any of the following characters: [*/\"\\ ;,.?<>:@#$%^&()=] and should consist of no more than 255 characters. • Tracking the link state of member ports in a LAG or MLAG is not supported	

link state tracking vlan

link state tracking vlan <vlan-id>
no link state tracking vlan <vlan-id>

Creates a vlan link state tracking group. All vlan members are automatically added into this group.

Syntax Description	vlan-id	ID of VLAN whose link state to track
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Role	admin	
Example	`switch (config)# link state tracking vlan 100`	
Related Commands		
Note	The maximum number of tracking groups/vlans is 64	

show link state tracking

show link state tracking [group <group-name> | vlan <vlan-id>]

Displays link state tracking configuration.

Syntax Description	group	displays link state tracking per tracking group
	vlan	displays link state tracking per VLAN
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Role	admin	

Example

```
switch (config)# show link state tracking

-------------------------------------------------------------------------------------
Group              Port Type      Interface          Admin Status   Operational Status
-------------------------------------------------------------------------------------
Vlan 100           Upstream       Eth1/54            Enabled        Down
Vlan 100           Downstream     Eth1/1             Enabled        Down (by tracking)
Vlan 100           Unassigned     Eth1/2             Enabled        Up
Vlan 101           Upstream       Eth1/54            Enabled        Down
Vlan 101           Downstream     Eth1/1             Enabled        Down (by tracking)
Vlan 101           Unassigned     Eth1/2             Enabled        Up
group1             Downstream     Eth1/1             Enabled        Down (by tracking)
```

Related Commands	
Note	

5.10 QinQ

A QinQ VLAN tunnel enables a service provider (SP) to segregate the traffic of different customers in their infrastructure, while still giving the customer a full range of VLANs for their internal use by adding a second 802.1Q VLAN tag to an already tagged frame.

So let us assume for example that an SP exists which needs to offer L2 connectivity to two corporations, “X” and “Y”, that have campuses located in both “A”, “B”. All campuses run Ethernet LANs, and the customers intend to connect through the SP’s L2 VPN network so that their campuses are in the same LAN (L2 network). Hence, it would be desirable for “X”, “Y” to have a single LAN each in both “A”, “B” which could easily exceed the VLAN limit of 4096 of the 802.1Q specification.

5.10.1 QinQ Operation Modes

QinQ can be enabled on a port or according to predefined conditions.

C-VLAN is the VLAN tag assigned to the ingress traffic of a QinQ-enabled interface.
S-VLAN is the VLAN tag assigned to the egress traffic of a QinQ-enabled interface.

- ACL-mode: Adding and removing S-VLAN is determined by an ACL-dependent action
- Port-mode: All ingress traffic to a specific QinQ-enabled interface is tagged with an additional VLAN 802.1Q tag (also known as S-VLAN). The S-VLAN ID is equal to that interface’s PVID (access VLAN).

  The S-VLAN tag is added regardless of whether the traffic is tagged or untagged. Traffic coming out from this port, has the S-VLAN stripped from it.

5.10.2 Configuring QinQ

➢ ***To configure QinQ:***

Step 1. Create the C-VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # exit
```

Step 2. Enter the configuration mode of an Ethernet, LAG, or MLAG interface. Run:

```
switch (config) # interface port-channel 100
```

Step 3. Change the switchport mode of the interface to enable QinQ. Run:

```
switch (config interface port-channel 100) # switchport mode dot1q-tunnel
```

Step 4. Change its port VLAN ID (PVID). This configures the S-VLAN. Run:

```
switch (config interface port-channel 100) # switchport access vlan 200
```

Step 5. Verify the configuration. Run:

```
switch (config interface port-channel 100) # show interface port-channel 100

Po100
```

```
  Admin state: Enabled
  Operational state: Up
  Description: N\A
  Mac address: 00:00:00:00:00:00
    MTU: 1500 bytes(Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control: receive off send off
  Actual speed: 1 X 40 Gbps
  Width reduction mode: disabled
  Switchport mode: dot1q-tunnel
  QoS mode: uniform
  MAC learning mode: Enabled
  Last clearing of "show interface" counters : Never
  60 seconds ingress rate: 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate: 0 bits/sec, 0 bytes/sec, 0 packets/sec

Rx
  0 packets
  0 unicast packets
  0 multicast packets
  0 broadcast packets
  0 bytes
  0 error packets
  0 discard packets

Tx
  0 packets
  0 unicast packets
  0 multicast packets
  0 broadcast packets
  0 bytes
  0 discard packets
switch (config interface port-channel 100) #
```

Step 6. Verify the configuration. Run:

```
switch (config interface port-channel 100) # show interfaces switchport
Interface       Mode         Access vlan       Allowed vlans
-------------------------------------------------------------------------
Eth1/1          access       1
Eth1/2          access       1
Eth1/3          access       1
Eth1/4          access       1
Eth1/5          access       1
Eth1/6          access       1
...
Eth1/27         access       1
Eth1/33         access       1
Eth1/34         access       1
Eth1/35         access       1
Eth1/36         access       1
Po400           dot1q-tunnel 200
```

5.10.3 Commands

switchport dot1q-tunnel qos-mode

switchport dot1q-tunnel qos-mode {pipe | uniform}
no switchport dot1q-tunnel qos-mode

Assigns QoS to the service provider's traffic.
The no form of the command resets the parameter value to its default.

Syntax Description	pipe	Gives the service provider's traffic the same QoS as the customer's traffic
	uniform	Gives the service provider's traffic QoS 0
Default	pipe	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.4.3000	
Role	admin	
Example	`switch (config  1/1) # switchport dot1q-tunnel qos-mode uniform` `switch (config  1/1) #`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan vlan	
Note		

5.11 Access Control List

An Access Control List (ACL) is a list of permissions attached to an object, to filter or match switches packets. When the pattern is matched at the hardware lookup engine, a specified action (e.g. permit/deny) is applied. The rule fields represent flow characteristics such as source and destination addresses, protocol and VLAN ID.

ACL support currently allows actions of *permit* or *deny* rules, and supports only ingress direction. ACL search pattern can be taken from either L2 or L3 fields, e.g L2/L3 source and destination addresses, protocol, VLAN ID and priority or TCP port.

5.11.1 Configuring Access Control List

Access Control List (ACL) is configured by the user and is applied to a port once the ACL search engine matches search criteria with a received packet.

➢ ***To configure ACL:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Create a MAC / IPv4 ACL (access-list) entity.

```
switch (config) mac access-list mac-acl
switch (config mac access-list mac-acl) #
```

Step 4. Add a MAC / IP rules to the appropriate access-list.

```
switch (config mac access-list mac-acl)seq-number 10 deny 0a:0a:0a:0a:0a:0a mask
ff:ff:ff:ff:ff:ff any vlan 6 cos 2 protocol 80
switch (config mac access-list mac-acl) #
```

Step 5. Bind the created access-list to an interface (slot/port or port-channel).

```
switch (config)
switch (config) #  1/1
switch (config  1/1) # mac port access-group mac-acl
```

5.11.2 ACL Actions

An ACL action is a set of actions can be activated in case the packet hits the ACL rule.

➢ ***To modify the VLAN tag of the egress traffic as part of the ACL "permit" rule:***

Step 1. Create access-list action profile:

Step 1a. Create an action access-list profile using the command `access-list action <action-profile-name>`.

Step 1b. Add rule to map a VLAN using the command `vlan-map <vlan-id>` within the action profile configuration mode.

Step 1c. Add action on a rule to strip the VLAN from a packet using the command `vlan-pop` within the action profile configuration mode.

Step 1d. Add action on a rule to append a VLAN to a packet using the command `vlan-push` within the action profile configuration mode.

Step 2. Create an access-list and bind the action rule:

a.Create an access-list profile using the command `{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list`.

b.Add access list rule using the command `deny/permit (action <action profile name>)`.

Step 3. Bind the access-list to an interface using the command `{ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group`.

```
Create an action profile and add vlan mapping action:
switch (config)# access-list action my-action
switch (config access-list action my-action)# vlan-map 20
switch (config access-list action my-action)# exit

Create an access list and bind rules:
switch (config)# mac access-list my-list
switch (config mac access-list my-list)# permit any any action my-action
switch (config mac access-list my-list)# exit

Bind an access-list to a port:
switch (config)#  1/1
switch (config  1/1)# mac access-list my-list
```

5.11.3 ACL Logging

A strong insight into the system is given by ACL logging. ACLs can log packets that pass through the switch, so the flows can later be analyzed.

A packet that hits an ACL with a log clause is passed to the logger. The logger writes the partial header of the packet (L2 or L3) to the syslog, with a timestamp and some additional information such as ingress interface and the VLAN to which the packet belongs.

To protect the system memory, a limited number of flows are collected for each time interval. If the number of flows for a specific time interval is exceeded, then no packets are logged for this time interval.

To further protect the system, a rate-limiter controls the number of packets passed to the CPU.

Only packets traversing the switch are logged. Packets that are passed to the CPU are not.

5.11.4 ACL Capability Summary

Table 53 summarizes the ACL capabilities supported by Onyx.

Table 53 - Summary of ACL Capability

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)	Scale
MAC	Permit Deny Remark[1]	N/A	DST MAC (with mask) SRC MAC (with mask) Protocol CoS VLAN-ID VLAN-group	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer	L2 port LAG MLAG RIF VLAN interface	18K
IPv4	Permit Deny Remark[1]	IP	DST IP (incl. subnets) SRC IP (incl. subnets)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer	L2 port LAG MLAG RIF VLAN interface	9K
		TCP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow			
		UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)			
		TCP-UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)			
		ICMP	DST IP (incl. subnets) SRC IP (incl. subnets) Code Type			

Table 53 - Summary of ACL Capability (Continued)

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)	Scale
IPv6	Permit Deny Remark[1]	IPv6	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer	L2 port LAG MLAG RIF VLAN interface	6K
		TCP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow			
		UDP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)			
		TCP-UDP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)			
		ICMPv6	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) Code Type			
MAC-UDK	Permit Deny Remark[1]	N/A	DST MAC (with mask) SRC MAC (with mask) Protocol CoS VLAN-ID VLAN-group UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer	L2 port LAG MLAG RIF VLAN interface	9K

Table 53 - Summary of ACL Capability (Continued)

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)	Scale
IPv4-UDK	Permit Deny Remark[1]	IP	DST IP (incl. subnets) SRC IP (incl. subnets) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer	L2 port LAG MLAG RIF VLAN interface	6K
		TCP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)			
		UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)			
		TCP-UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)			
		ICMP	DST IP (incl. subnets) SRC IP (incl. subnets) Code Type UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)			

1. No keys & actions

5.11.5 Commands

{ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list

{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list <acl-name>
no {ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list <acl-name>

Creates an ACL table and enters its configuration mode.
The no form of the command deletes the ACL table.

Syntax Description	ipv4 | mac	IPv4 or MAC – access list
	acl-name	User-defined string for the ACL
Default	No ACL available by default.	
Configuration Mode	config	
History	3.1.1400	
	3.6.5000	Added ipv6, ipv4-udk, and mac-udk parameters
Role	admin	
Example	`switch (config)# mac access-list my-mac-list` `switch (config mac access-list my-mac-list)#`	
Related Commands	ipv4/port access-group	
Note	• Each table has its own set of predefine keys • The mac-udk and ipv4-udk options add an extra UDK to the standard MAC and IPv4 tables • When a new access-list is created, its default bind port is L2 port	

bind-point rif

	bind-point rif **no bind-point rif** Changes the ACL table bind point from L2 port mode to L3 port. The no form of the command resets this parameter to its default.
Syntax Description	N/A
Default	L2 port
Configuration Mode	config mac access-list
History	3.6.5000
Role	admin
Example	`switch (config mac access-list my-mac-list)# bind-point rif`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list
Note	• The bind point may only be changed when an ACL table is empty (no rules) and unbound • This command is used to attach ACLs to interface VLANs only

remark

	[<seq-number>] remark <string> **no [<seq-number>] remark <string>** Creates a remark rule from an ACL table. The no form of the command deletes a remark rule from an ACL table.
Syntax Description	N/A
Default	N/A
Configuration Mode	config mac access-list
History	3.6.5000
Role	admin
Example	`switch (config mac access-list my-mac-list)# remark “1st group”`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list
Note	• The remark rule has a sequence number like standard rules and it can be displayed when showing all rules of ACL table • This rule has no effect on traffic and it is only for management purposes

shared-counter

shared-counter <counter-name>
no shared-counter <counter-name>

Creates a shared counter.
The no form of the command deletes a shared counter.

Syntax Description	counter-name	Shared counter name
Default	N/A	
Configuration Mode	config mac access-list	
History	3.6.5000	
Role	admin	
Example	`switch (config mac access-list my-mac-list)# shared-counter myCounter`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list	
Note	• When creating a new shared counter, it is created only in the scope of the ACL table it has been initially created on and cannot be shared across multiple ACL tables • A shared counter cannot be deleted when attached to rules	

clear shared-counters

clear shared-counters [<counter-name>]

Resets all shared counters in ACL table or a specific shared counter.

Syntax Description	counter-name	Shared counter name
Default	N/A	
Configuration Mode	config mac access-list	
History	3.6.5000	
Role	admin	
Example	`switch (config mac access-list my-mac-list)# clear shared-counters`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter	
Note		

clear counters

clear counters [<seq-number>]

Resets all counters (including shared counters) in ACL table or a specific counter.

Syntax Description	seq-number	The sequence number of the rule whose counter to reset
Default	N/A	
Configuration Mode	config mac access-list	
History	3.6.5000	
Role	admin	
Example	`switch (config mac access-list my-mac-list)# clear counters 10`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter	
Note		

{ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list clear counters

	{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list clear counters Resets all counters (including shared counters) on all ACL tables of the same type.
Syntax Description	N/A
Default	N/A
Configuration Mode	config mac access-list
History	3.6.5000
Role	admin
Example	`switch (config)# ipv4 access-list clear counters`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter
Note	

{ipv4/ipv6/mac/ipv4-udk/mac-udk} port access-group

{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} port access-group <acl-name>
no {ipv4 | ipv6 | mac | ipv4-udk | mac-udk} port access-group <acl-name>

Binds an ACL to the interface.
The no form of the command unbinds the ACL from the interface.

Syntax Description	ipv4 | mac	IPv4 or MAC – access list
	acl-name	ACL name
Default	No ACL is bind by default.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
	3.3.4500	Added MPO configuration mode
	3.6.5000	Added new parameters
Role	admin	
Example	`switch (config  1/1) # mac port access-group my-list` `switch (config  1/1) #`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list	
Note	The access control list should be defined prior to the binding action.	

deny/permit (MAC ACL rule)

[seq-number <sequence-number>] {permit | deny} ip {<source-mac> mask <mac_mask> | any} {<dest-mac> mask <mac_mask> | any} [protocol <protocol_num>] [cos <cos>] [vlan <vlan_id>] [vlan-mask <vlan_mask>] [action <action-name>] [log] [counter | shared-counter <name>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for MAC ACL.
The no form of the command deletes a rule from the MAC ACL.

Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule. Range is:1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-mac> mask <mac_mask> | any	Sets source MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the source MAC.
	<dest-mac> mask <mac_mask> | any	Sets destination MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the destination MAC.
	protocol	Sets the Ethertype field value from the MAC address. Range: 0x0000-0xffff.
	cos	Sets the COS (priority bit) field. Range is: 0-7.
	vlan <vlan_id>	Sets the VLAN ID field. Range is 1-4094.
	vlan-mask <vlan-mask>	Sets VLAN group. Range: 0x0000-0x0FFF.
	action	Action name (free string)
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).

	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config mac acl	
History	3.1.1400	
	3.3.4500	Added vlan-mask parameter
	3.5.1000	Updated seq-number parameter
	3.6.5000	Added log, counter, and shared-counter parameters
	3.6.6000	Added policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config mac access-list my-list) # seq-number 10 deny 0a:0a:0a:0a:0a:0a mask ff:ff:ff:ff:ff:ff any vlan 6 cos 2 protocol 80` `switch (config mac access-list my-list) #`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Note	• VLAN and VLAN group cannot be used in the same command • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv4 ACL rule)

[seq-number <sequence-number>] {permit | deny} ip {<source-ip> mask <ip> | [any]} {<dest-ip> mask <ip> | [any]} [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 ACL.
The no form of the command deletes a rule from the IPv4 ACL.

Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule. Range is: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	{any | <source-ip> mask <ip>}	Sets source IP and optionally sets a mask for that IP address. The “any” option causes the rule to not check the source IP. Range: 0-255.
	{any | <destination-ip> mask <ip>}	Sets destination IP and optionally sets a mask for that IP. The “any” option causes the rule to not check the destination IP.
	action	Action needs to be defined before attaching to rule
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000.

	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.3.4302	Updated syntax description of mask <ip> parameter
	3.5.1000	Updated seq-number parameter
	3.6.5000	Added log, counter, and shared-counter parameters
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list) # deny ip any any action act shared-counter`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Note	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv4 TCP ACL rule)

[seq-number <sequence-number>] {deny | permit} tcp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 TCP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range is: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	eq-source <src-port>	TCP source port number. Range: 0-65535.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	eq-destination <dest-port>	TCP destination port number. Range: 0-65535.
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	established	Matches flows which are in established state (“ack” or “rst” flags are set)

	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, policer, and extra flag parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list)# permit tcp any any src-port 200 dest-port-range 200 400 established` `switch (config ipv4 access-list my-list)# permit tcp any any ns 0 policer packets rate 1 k burst 2050`	

Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	• L4 ports are valid • It is possible to attach the rule to a unique policer, or to create a policer only for the rule

deny/permit (IPv4 TCP-UDP/UDP ACL rule)

[seq-number <sequence-number>] {deny | permit} {tcp-udp | udp} {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 TCP-UDP/UDP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	eq-source <src-port>	TCP-UDP/UDP source port number. Range: 0-65535.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	eq-destination <dest-port>	TCP-UDP/UDP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule

	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list)# permit tcp-udp any any eq-destination 100 eq-source 300` `switch (config ipv4 access-list my-list)# permit udp any any eq-destination 100 eq-source 300`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv4 ICMP ACL rule)

[seq-number <sequence-number>] {deny | permit} icmp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [eq-code <icmp-code>] [eq-type <icmp-type>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 ICMP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value. Range: 0-255.
	eq-type	Matches ICMP type value. Range: 0-255.
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).

	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.2002	Added ICMP parameters
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list)# permit icmp any any eq-code 10 eq-type 155`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• ICMP code must be specified in conjunction with an ICMP type. If ICMP type is specified but no ICMP code is specified, the rule matches all ICMP packets of the given type • If no ICMP type or code are specified, the rule matches all ICMP packets from the specified source/destination address • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv6 ACL rule)

[seq-number <sequence-number>] {permit | deny} ip {<src-ipv6>/<mask-len> | any} {<dest-ipv6>/<mask-len> | any} [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates an IPv6 ACL rule with a specific protocol.
The no form of the command deletes a rule from the IPv6 ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<src-ipv6>/<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The parameter “any” ignores the source IP.
	<dest-ipv6>/<mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The parameter “any” ignores the destination IP.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7

	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv6 access-list my-list) # permit ip 2:2::/32 any` `switch (config ipv6 access-list my-list) # permit ip any any policer name`	
Related Commands		
Note	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • The fields eq-code (icmp-code) and eq-type (eq-type) are valid only for ICMP rules • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv6 TCP ACL rule)

[seq-number <sequence-number>] {permit | deny} tcp {<source-ipv6> /<mask-len> | any} {<dest-ipv6> /<mask-len> | any} [src-port <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | dest-port-range <from> <to>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates an IPv6 ACL rule with a specific protocol.
The no form of the command deletes a rule from the IPv6 ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ipv6> /<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> /<mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	established	Matches flows which are in established state (“ack” or “rst” flags are set)
	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option

	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, policer, and flag parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv6 access-list my-list) # permit tcp any 10:10:12::/48`	
Related Commands		
Note	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv6 TCP-UDP/UDP ACL rule)

[seq-number <sequence-number>] {permit | deny} {tcp-udp | udp} {<source-ipv6> /<mask-len> | any} {<dest-ipv6> /<mask-len> | any} [src-port <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | dest-port-range <from> <to>] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates an IPv6 ACL rule with a specific protocol.
The no form of the command deletes a rule from the IPv6 ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ipv6> /<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> /<mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.

	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv6 access-list my-list) # permit udp 2:2::/32 10:10:12::/48`	
Related Commands		
Note	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv6 ICMPv6 ACL rule)

[seq-number <sequence-number>] {permit | deny} icmpv6 {<source-ipv6> / <mask-len> | any} {<dest-ipv6> /<mask-len> | any} [code <icmp-code>] [type <icmp-type>] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates an IPv6 ACL rule with a specific protocol.
The no form of the command deletes a rule from the IPv6 ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ipv6> /<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> /<mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value Range: 0-255
	eq-type	Matches ICMP type value Range: 0-255
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer

	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv6 access-list my-list) # permit icmpv6 any any eq-code 10 eq-type 155`	
Related Commands		
Note	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (MAC UDK ACL rule)

[seq-number <sequence-number>] {deny | permit} {<source-mac> mask <mac-mask> | any} {<dest-mac> mask <mac-mask> | any} [protocol <protocol-num>] [cos <cos>] [vlan <vlan-id>] [vlan-mask <vlan_mask>] [action <action-name>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a MAC-UDK ACL rule.
The no form of the command deletes a rule from MAC UDK ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range:1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-mac> mask <mac-mask> | any	Sets source MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the source MAC.
	<dest-mac> mask <mac-mask> | any	Sets destination MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the destination MAC.
	protocol	Sets the Ethertype filed value from the MAC address Range: 0x0000-0xffff
	cos	Sets the COS (priority bit) field. Range: 0-7.
	vlan <vlan-id>	Sets the VLAN ID field. Range: 1-4094.
	vlan-mask <vlan-mask>	Sets VLAN group. Range: 0x0000-0x0FFF.
	action	Action name (free string)
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	policer	Attaches shared policer to a rule

	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config mac-udk acl	
History	3.6.5000	
	3.6.6000	Added policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config mac-udk access-list mac_udk_acl) # permit any any udk myUdk 10 mask 0xff`	
Related Commands		
Note	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • UDK fields must come at the end of the rule configuration • The default mask is 0xff–0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • Values and masks of the UDK can be decimal or hexadecimal • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	
	•	

deny/permit (IPv4 UDK ACL rule)

[seq-number <sequence-number>] {permit | deny} ip {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 ACL.
The no form of the command deletes a rule from the IPv4 ACL.

Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule. Range is: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	{any | <source-ip> mask <ip>}	Sets source IP and optionally sets a mask for that IP address. The “any” option causes the rule to not check the source IP. Range: 0-255.
	{any | <destination-ip> mask <ip>}	Sets destination IP and optionally sets a mask for that IP. The “any” option causes the rule to not check the destination IP.
	action	Action needs to be defined before attaching to rule
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer

	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list) # deny ip any any action act shared-counter`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Note	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • UDK fields must come at the end of the rule configuration • The default mask is 0xff–0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • Values and masks of the UDK can be decimal or hexadecimal • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

deny/permit (IPv4 TCP UDK ACL rule)

[seq-number <sequence-number>] {deny | permit} tcp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 TCP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range is: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> [mask <ip>] | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> [mask <ip>] | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	eq-source <src-port>	TCP source port number. Range: 0-65535.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	eq-destination <dest-port>	TCP destination port number. Range: 0-65535.
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule

	established	Matches flows which are in established state (“ack” or “rst” flags are set)
	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, policer, and flag parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	

Example	`switch (config ipv4 access-list my-list)# permit tcp any any src-port 200 dest-port-range 200 400 established`
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	• UDK fields must come at the end of the rule configuration • The default mask is 0xff–0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule

deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)

[seq-number <sequence-number>] {deny | permit} {tcp-udp | udp} {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 TCP-UDP/UDP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port. Note: User may only choose one of the following options to configure source port: src-port; eq-source.
	eq-source <src-port>	TCP-UDP/UDP source port number. Range: 0-65535.
	src-port-range	Sets a range of L4 source ports to match. Note: User may configure either a single source port or a range.
	dest-port	L4 destination port. Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination.
	eq-destination <dest-port>	TCP-UDP/UDP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option

counter	Attaches a unique counter to rule
shared-counter	Attaches a predefined shared-counter to rule
udk	UDK name must be set by user before the rule configuration
val	The value of the UDK (up to 4 bytes)
mask	Mask for the UDK value
ecn	ECN ACL filter. Value: 0-3.
ttl	Time to live ACL filter. Value: 0-225.
dscp	DSCP ACL filter. Value: 0-63.
policer	Attaches shared policer to a rule
bytes	Attaches bytes type policer
bits	Attaches bits type policer
packets	Attaches packets type policer
rate	Policer rate value: 100-1000000000000
k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7

Default	No rule is added by default to access control list Default sequence number is by increments of 10
Configuration Mode	config ipv4 acl
History	3.6.5000
	3.6.6000 — Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx — Added bits, switch-priority and tc parameters
Role	admin
Example	`switch (config ipv4 access-list my-list)# permit tcp-udp any any eq-destination 100 eq-source 300` `switch (config ipv4 access-list my-list)# permit udp any any eq-destination 100 eq-source 300`

Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	• UDK fields must come at the end of the rule configuration • The default mask is 0xff–0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule

deny/permit (IPv4 ICMP UDK ACL rule)

[seq-number <sequence-number>] {deny | permit} icmp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [eq-code <icmp-code>] [eq-type <icmp-type>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]}
no <sequence-number>

Creates a rule for IPv4 ICMP ACL.
The no form of the command deletes a rule from the ACL.

Syntax Description		
	sequence-number	Optional parameter to set a specific sequence number for the rule. Range: 1-65535.
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value. Range: 0-255.
	eq-type	Matches ICMP type value. Range: 0-255.
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer

	bits	Attaches bits type policer
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo (10^3), mega (10^6), or giga (10^9).
	burst	Sets burst to policer. Max size: 64. If no burst is configured, the default value for type "packets" is 100 and for "bytes" is 10000.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.00xx	Added bits, switch-priority and tc parameters
Role	admin	
Example	`switch (config ipv4 access-list my-list)# permit icmp any any eq-code 10 eq-type 155`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• ICMP code must be specified in conjunction with an ICMP type. If ICMP type is specified but no ICMP code is specified, the rule matches all ICMP packets of the given type. • If no ICMP type or code are specified, the rule matches all ICMP packets from the specified source/destination address. • UDK fields must come at the end of the rule configuration • The default mask is 0xff–0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule	

port access-group (IPv4/IPv4 UDK/IPv6/MAC/MAC UDK)

{ipv4 | ipv4-udk | ipv6 | mac | mac-udk} port access-group <acl-name>
no {mac | ipv4 | ipv6 | mac-udk | ipv4-udk} port access-group

Attaches an ACL table with bind-point RIF to a VLAN interface.
The no form of the command unmaps ACL table with bind-point RIF from a VLAN interface.

Syntax Description	acl-name	ACL table name
Default	N/A	
Configuration Mode	config interface vlan	
History	3.6.5000	
Role	admin	
Example	`switch (config interface vlan 10)# ipv4 port access-group ipv4_acl2`	
Related Commands	show access list summary	
Note	• Only ACL tables with bind-point set to RIF can be attached to a VLAN interface • Interface VLAN must be configured before binding operation	

access-list action

access-list action <action-profile-name>
no access-list action <action-profile-name>

Creates access-list action profile and entering the action profile configuration mode. The no form of the command deletes the action profile.

Syntax Description	action-profile-name	Given name for the profile.
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
Role	admin	
Example	`switch (config)# access-list action my-action` `switch (config access-list action my-action)#`	
Related Commands		
Note		

access-list log

access-list log [interval <int_num>] [memory <packet_num>] [syslog <packet_-num>]
no access-list log [interval <int_num>] [memory <packet_num>] [syslog <packet_num>]

Configures access list logger.
The no form of the command resets parameters for access list logger.

Syntax Description	interval	Logging interval length in minutes Range: 1min-24hrs
	memory	Maximal number of packets to save in memory Range: 1-3600
	syslog	Maximal number of packets to show in syslog Range: 1-3600
Default	N/A	
Configuration Mode	config	
History	3.6.5000	
Role	admin	

Example	`switch (config)# access-list log interval 10`
	`switch (config)# access-list log memory 300`
	`switch (config)# access-list log syslog 200`	

Related Commands		
Note	• The packet number in syslog configuration must not be greater than the maximal packets number in memory • When configuring interval, the interval will restart resulting in a log dump to syslog and memory clear	

vlan-map

vlan-map <vlan-id>
no vlan-map

Adds action to map a new VLAN to the packet (in the ingress port or VLAN).
The no form of the command removes the action to map a new VLAN.

Syntax Description	vlan-id 1-4094.
Default	N/A
Configuration Mode	config acl action
History	3.2.0230
Role	admin
Example	`switch (config access-list action my-action)# vlan-map 10`
Related Commands	
Note	

vlan-pop

vlan-pop

Pops VLAN frames from traffic.

Syntax Description	vlan-id	VLAN ID: 1-4094.
Default	N/A	
Configuration Mode	config acl action	
History	3.4.3000	
Role	admin	
Example	`switch (config access-list action my-action)# vlan-pop`	
Related Commands		
Note		

vlan-push

vlan-push <vlan-id>

Pushes (or adds) VLAN frames to traffic.

Syntax Description	vlan-id	VLAN ID: 1-4094
Default	N/A	
Configuration Mode	config acl action	
History	3.4.3000	
Role	admin	
Example	`switch (config access-list action my-action)# vlan-push 10`	
Related Commands		
Note		

show ipv4 access-lists

show ipv4 access-lists <access-list-name>

Displays configuration of IPv4 rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.3.4500	Updated Example
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show ipv4 access-lists my-list

Table Type: ipv4
Table Name: my-list
Bind-point: port
```

seq-number	p/d	protocol	s-ipv4	d-ipv4	sport/type	end-sport	dport/code	end-dport	tcp-control	action	counter	Packets	ttl	ecn	dscp	policer	log
10	permit	ip	any	any	any	none	any	none	N/A	none	N/A	N/A	none	none	none	none	NO
20	permit	ip	any	any	any	none	any	none	N/A	none	N/A	N/A	none	none	none	YES	NO

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show ipv4-udk access-lists

show ipv4-udk access-lists <access-list-name>

Displays configuration of IPv4 UDK rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show ipv4-udk access-lists my-list

Table Type: ipv4-udk
Table Name: my-list
Bind-point: port
```

seq-number	p/d	protocol	s-ipv4	d-ipv4	sport/type	end-sport	dport/code	end-dport	tcp-control	action	counter	Packets	udk	ttl	ecn	dscp	policer	log
7	permit	tcp	any	any	any	none	any	none	any	none	N/A	N/A		none	none	none	none	NO
8	deny	tcp	1.1.1.1/32	any	any	none	any	none	-U +F	none	N/A	N/A	aaa value 5	none	none	none	none	NO
10	permit	tcp	1.1.1.1/32	2.2.2.2/32	any	none	any	none	+P-R	none	N/A	N/A	bbb value 6 mask 0x8	none	none	none	none	NO

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show ipv6 access-lists

show ipv6 access-lists <access-list-name>

Displays configuration of IPv6 rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show ipv6 access-lists my-list

Table Type: ipv6
Table Name: my-list
Bind-point: port
```

seq-number	p/d	protocol	s-ipv6	d-ipv6	sport/type	end-sport	dport/code	end-dport	tcp-control	action	counter	Packets	ttl	ecn	dscp	policer	log
10	permit	ip	any	any	any	none	any	none	N/A	none	N/A	N/A	33	none	none	none	YES
20	permit	ip	any	any	any	none	any	none	N/A	none	N/A	N/A	none	none	none	none	NO
30	permit	ip	any	any	any	none	any	none	N/A	none	N/A	N/A	none	none	none	none	NO

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show mac access-lists

show mac access-lists <access-list-name>

Displays configuration of MAC rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.3.4500	Updated Example
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show mac access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------------------------------
seq-number  p/d     smac              dmac              protocol  cos    vlan   vlan-mask  action  counter     Packets       policer     log
-------------------------------------------------------------------------------------------------------------------------------------
10          permit  any               any               any       any    any    N/A        none    N/A         N/A           roe         NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show mac access-lists summary

show mac access-lists <access-list-name>

Displays configuration of MAC rules in a specific table.

Syntax Description	access-list-name ACL name
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Role	admin
Example	

```
switch (config) # show mac access-lists summary

-------------------------------------------------------------------------------------
Table type    Table Name       Bind Point    Total entries    Bound to interfaces
-------------------------------------------------------------------------------------
mac           mac1             port          1                Eth1/16
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show mac-udk access-lists

show mac-udk access-lists <access-list-name>

Displays configuration of MAC UDK rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show mac-udk access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port
```

seq-number	p/d	smac	dmac	protocol	cos	vlan	vlan-mask	action	counter	Packets	udk	policer	log
10	permit	any	any	any	any	any	N/A	none	N/A	0		YES	NO
20	permit	any	any	any	any	any	N/A	none	N/A	N/A		none	NO

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show access-lists action

show access-lists action <action-profile-name>

Displays the access-list action profiles summary.

Syntax Description	action-profile-name	Filter the table according to the action profile name.
	summary	Display summary of the action list.
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.7.1000	
Role	admin	
Example	witch (config)# show access-lists action my-action Access-list Action my-action ========== Mapped_Vlan_ID |Mapped_port |Counter_set |Policer_ID | ========== 10 |N/A |N/A |N/A | switch (config)#	
Related Commands		
Note		

show mac-udk access-lists

show mac-udk access-lists <access-list-name>

Displays configuration of MAC UDK rules in a specific table.

Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config) # show mac-udk access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------------------------------------
seq-number  p/d     smac            dmac            protocol  cos   vlan   vlan-mask  action  counter     Packets      udk                      policer     log
-------------------------------------------------------------------------------------------------------------------------------------------
10          permit  any             any             any       any   any    N/A        none    N/A         0                                    YES         NO
20          permit  any             any             any       any   any    N/A        none    N/A         N/A                                  none        NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Note	

show access-lists log config

show access-lists log config <action-profile-name>

Displays the access-list log configuration information.

Syntax Description	action-profile-name	Filter the table according to the action profile name.
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.6.8008	Updated Example
Role	admin	

Example	`witch (config)# show access-lists log config`
	
	`access-list log configuration:`
	`  Memory packets    : 1000`
	`  Syslog packets    : 10`
	`  Interval (minutes): 1`	

Related Commands		
Note		

show access-lists policers (ipv4/ipv4-udk/ipv6/mac/mac-udk)

show {ipv4 | ipv4-udk | ipv6 | mac | mac-udk} access-lists <access-list-name> policers [name | seq-number]

Displays all configured policers on a specific ACL table.

Syntax Description	access-list-name	ACL name
	name	Policer name filter
	seq-number	Filter by sequence number
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	
Example		

```
switch (config) # show ipv6 access-lists my-list policers

-------------------------------------------------------------------------------------
Name            Type     Rate           Burst     Sequence Number
-------------------------------------------------------------------------------------
pol             packets  1000           200       50,60,70
rom             packets  1000           200       80
N/A             bytes    12345          20000     40
```

Related Commands	
Note	

show access-lists shared-counters (ipv4/ipv4-udk/ipv6/mac/mac-udk)

show {ipv4 | ipv4-udk | ipv6 | mac | mac-udk} access-lists <access-list-name> shared-counters

Displays all configured shared-counters on a specific ACL table.

Syntax Description	access-list-name ACL name
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin

Example	`switch (config mac access-list my-list) # show mac access-lists mac_acl shared-counters`
	`-----------------------------------------------`
	`counter packets total Rules rule IDs`
	`-----------------------------------------------`
	`cnt1 0 3 20 30 40`
	`cnt2 0 2 50 60`
	`cnt3 0 1 70`

Related Commands	
Note	• For each configured shared counter it also displays the counter value (packets), the number of rules attached to this counter and the rule IDs • Up to 5 rule IDs are displayed even though there is no limitation on how many rules can be attached to a counter

show access-lists summary

	show [ipv4 | mac | ipv6 | ipv4-udk | mac-udk] access-lists summary Displays the summary of number of rules per ACL, and the interfaces attached.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.6.5000 Updated Example
Role	admin

Example	`switch (config) # show access-lists summary`
	`---------------------------------------------------------------------------------`
	`Table type  Table Name  Bind type  Total entries  Bound to interfaces`
	`---------------------------------------------------------------------------------`
	`mac  aaa  port  0  Mpo55`
	`ipv4  ddd  port  1  Eth1/3, Po1`
	`ipv4  ggg  rif  0  VlanIf555`
	`ipv6  table1  port  9  Eth1/9`

Related Commands	
Note	

show access-lists log

show access-lists log [last <num>]

Displays captured packets on all access list rules.

Syntax Description	num Number of packets to show
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin

Example	`switch (config) # show access-lists log`
	`Log status: Normal`
	
	`Log MAC rules:`
	`----------`
	`IF Table(rule) Source MAC Dest MAC Ethertype VLAN Hits`
	`----------`
	`1/2 mac_al_log(10) 44:44:44:44:44:44 22:22:22:22:22:22 IPv4 N/A 5`
	
	`Log IPv4 rules:`
	`----------`
	`IF Table(rule) Source IPv4 Dest IPv4 Protocol Source port Dest port Hits`
	`----------`
	`1/3 ipv4_al_lo(10) 1.1.1.1 2.2.2.2 UDP 44 33 11`

Related Commands	
Note	

show access-lists log config

	show access-lists log config Displays configuration of access-list logger.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin

Example	`switch (config) # show access-lists log config`
	`access-list log configuration:`
	`  Memory packets:  1000`
	`  Syslog packets:  10`
	`  Interval (minutes): 60`

Related Commands	
Note	

5.12 OpenFlow

Onyx supports OpenFlow 1.3. OpenFlow is a network protocol that facilitates direct communication between network systems via Ethernet. Software Defined Networks (SDN) allows a centralist management of network equipment. OpenFlow allows the SDN controller to manage SDN equipment. The OpenFlow protocol allows communication between the OpenFlow controller and OpenFlow agent.

OpenFlow is useful to manage switches and allow applications running on the OpenFlow controller to have access to the switch's data path and provide functionality such as flow steering, security enhancement, traffic monitoring and more.

The OpenFlow controller communicates with the OpenFlow switch over secured channel using OpenFlow protocol.

An OpenFlow switch contains a flow table which contains flows inserted by the OpenFlow controller. And the OpenFlow switch performs packet lookup and forwarding according to those rules.

Mellanox OpenFlow switch implementation is based on the hybrid model, allowing the coexistence of an OpenFlow pipeline and a normal pipeline. In this model, a packet is forwarded according to OpenFlow configuration, if such configuration is matched with the packet parameters. Otherwise, the packet is handled by the normal (regular forwarding/routing) pipeline.

The OpenFlow specification defines:

> "OpenFlow-hybrid switches support both OpenFlow operation and normal Ethernet switching operation, i.e. traditional L2 Ethernet switching, VLAN isolation, L3 routing (IPv4 routing, IPv6 routing...), ACL and QoS processing. Those switches must provide a classification mechanism outside of OpenFlow that routes traffic to either the OpenFlow pipeline or the normal pipeline. For example, a switch may use the VLAN tag or input port of the packet to decide whether to process the packet using one pipeline or the other, or it may direct all packets to the OpenFlow pipeline."

Utilizing the built-in capabilities of the hybrid switch/router is the main benefit of the hybrid mode. It increases network performance and efficiency – faster processing of new flows as well as lower load on the controllers. The hybrid switch processes non-OpenFlow data through its local management plane and achieve better efficiency and use of resources, compared to the pure OpenFlow switch.

5.12.1 Flow Table

The flow table contains flows which are used to perform packet lookup, modification and forwarding. Each flow has a 12 tuple key. The key is used in order to classify a packet into a certain flow. The key contains the flowing fields: ingress port, source MAC, destination MAC, EtherType, VLAN ID, PCP, source IP, destination IP, IP protocol, IP ToS bits, TCP/UDP source port and TCP/UDP destination port.

The flow key can have a specific value for each field or wildcard which signals to the switch to ignore this part of the key.

Each packet passes through the flow table once a match is found; the switch performs the actions configured to the specific flow by the OpenFlow controller.

Up-keeping a flow table enables the switch to forward incoming traffic with a simple lookup on its flow table entries. OpenFlow switches perform a check for matching entries on, or ignore using a wildcard, specific fields of the ingress traffic. If the entry exists, the switch performs the action associated with that flow entry. Packets without a flow entry match are forwarded according to the normal pipeline (hybrid switch).

Every flow entry contains one of the following parameters:

1. Header fields for matching purposes with each entry containing a specific value or a wildcard which could match all entries.
2. Matching packet counters which are useful for statistical purposes, in order to keep track of the number of packets.
3. Actions which specify the manner in which to handle the packets of a flow which can be any of the following:
   - Forwarding the packet
   - Dropping the packet
   - Forwarding the packet to the OpenFlow controller
   - Modifying the VLAN, VLAN priority (PCP), and/or stripping the VLAN header

5.12.2 OpenFlow 1.3 Work Flow

The OpenFlow (OF) pipeline is deployed in parallel to the usual Onyx pipeline.

The ingress port must is deployed in hybrid mode so as to serve both the OF and normal Onyx pipeline.

Figure 23: Onyx and OpenFlow Pipeline

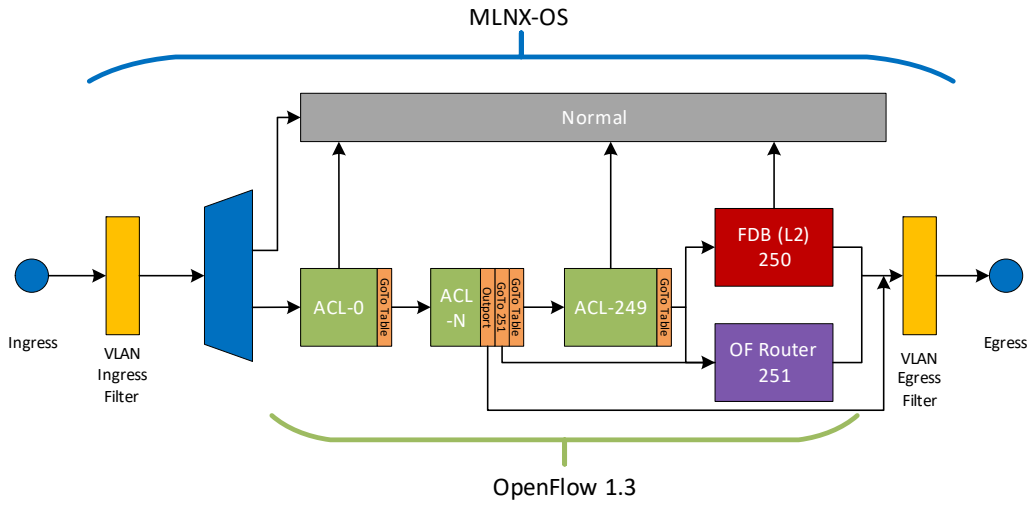

The ingress packet which passes the VLAN filter and is a match to the user ACL tables either progresses to the regular Onyx flow, or the OpenFlow pipeline depending on the port coupling.

Table 54 presents a general summary of the capabilities of the OpenFlow 1.3 pipeline, which are described in detail further on in the document.

Table 54 - OpenFlow 1.3 Pipeline Capabilities Summary Table

Table	Match	Actions	Group	Meters

ACLs [0-249]	• in_port
	• dl_src
	• dl_dst
	• dl_type
	• vlan_vid
	• vlan_pcp
	• ip_src
	• ip_dst
	• ipv6_dst
	• ipv6_src
	• ip_proto
	• ip_dscp
	• ip_ecn
	• ip_ttl
	• 14_src_poert
	• 14_dst_port
	• tunnel_id
	• metadata 0xFFF
	• mpls_label
	• Table must be configured using "openflow table match-keys" to support the following fields:
	  • ip_src_inner
	  • ip_dst_inner
	  • ignr_eth_type
	(Dynamic key)
	(Arbitrary mask)
		• Push/pop VLAN
		• SET_TTL
		• DEC_TTL
		• goto_table
		• Set queue
		• Eth SRC/DST MAC
		• VLAN ID
		• PCP
		• DSCP
		• ECN
		• Output
		• Group
		• Meters
		• Normal
			• ALL – Output ports
			• Select – {weights} Output ports (without LAG)
			• FF – Output ports
				• KBps/ PKTs – {Burst}
				• Drop

FDB [250]	• vlan_vid • dl_dst **(Exact match)**	• OUTPUT • DROP • Normal	Select – {Weights} Output ports (without LAG)	N/A
Router [251]	• ipv4_dst • ipv6_dst **(LPM)**	• DEC_TTL • SET_DMAC • OUTPUT • DROP **(Must have DEC_TTL and SET_DMAC when output action is implemented)**	Select – {Weights} output ports + set_dmac + dec_tll	N/A

5.12.2.1 ACL Rule Tables (0-249)

An Access Control List (ACL) is a list of permissions attached to an object, to filter or match switches packets. When the pattern is matched at the hardware lookup engine, a specified action (e.g. permit/deny) is applied. The rule fields represent flow characteristics such as source and destination addresses, protocol and VLAN ID.

ACL support currently allows actions of permit or deny rules, and supports only ingress direction. ACL search pattern can be taken from either L2 or L3 fields.

5.12.2.1.1 Supported ACL Matching Rules

Ingress packets, arriving the ACL, are matched against any combination of the following parameters (defined as the key):

- OXM_OF_METADATA – matches according to metadata
- OXM_OF_IN_PORT – matches according to ingress port (exact match or wildcard)
- OXM_OF_ETH_SRC – matches source MAC address
- OXM_OF_ETH_DST – matches destination MAC address
- OXM_OF_ETH_TYPE – matches EtherType

When match rule is set to match eth_type 9100, VLAN ID matching does not work.

- OXM_OF_VLAN_VID – matches VLAN ID
- OXM_OF_VLAN_PCP – matches priority level
- OXM_OF_IPV4_SRC – matches source IPv4 address
- OXM_OF_IPV4_DST – matches destination IPv4 address
- OXM_OF_IPV6_SRC – matches source IPv6 address
- OXM_OF_IPV6_DST – matches destination IPv6 address
- OXM_OF_IPV6_ND_TARGET

OXM_OF_IPV6_ND_TARGET match rule is not supported.

- OXM_OF_IP_PROTO – matches IP protocols (exact match or wildcard)
- OXM_OF_IP_DSCP – matches IP DSCP field (exact match or wildcard)
- OXM_OF_IP_ECN – matches network ECN (exact match or wildcard)
- OXM_OF_NW_TTL – matches network TTL (exact match or wildcard)
- OXM_OF_TCP_SRC – matches source TCP
- OXM_OF_TCP_DST – matches destination TCP
- OXM_OF_UDP_SRC – matches source UDP
- OXM_OF_UDP_DST – matches destination UDP

- OXM_OF_SCTP_SRC – matches source SCTP
- OXM_OF_SCTP_DST – matches destination SCTP
- OXM_OF_ICMPV4_TYPE – matches ICMP type
- OXM_OF_ICMPV4_CODE – matches ICMP code
- OXM_OF_ARP_OP – matches ARP OP code
- OXM_OF_ARP_SPA – matches sender protocol address
- OXM_OF_ARP_TPA – matches target protocol address

There is a default set of match keys configured. To see what it is, please run the command “show openflow table match-keys” on your machine. To alter it, please use the command “openflow table match-keys”.

5.12.2.1.2Non-standard Matches

OpenFlow 1.3 is able to match non-standard OpenFlow matching rules by mapping them to standard ones. The following non-standard matches are supported:

- Matching source/destination IPv4 address encapsulated with MPLS labels (up to 6 MPLS labels can be skipped)

  – ip_src_inner/ip_dst_inner is mapped to OXM_OF_IPV4_SRC, OXM_OF_IPV4_DST

Please see the example configuration below:

- Table configuration:

```
openflow table 0 match-keys dl_dst dl_src dl_type mpls_label vlan_vid
openflow table 10 match-keys ignr_eth_type ip_dst_inner ip_src_inner
```

  The ignr_eth_type is needed to ignore the Ethertype of IP that is required by OpenFlow to set to as a prerequisite to match on IP addresses

- Rules:

```
openflow add-flows 1 table=0,mpls,mpls_label:32,actions=goto_table=10
openflow add-flows 2 table=10,ip,nw_src=10.10.10.0/24,nw_dst=10.10.20.0/24,
actions=output:127
```

  The above matches IP address from 10.10.10.0/24 to 10.10.20.0/24 which have MPLS label 32 as the first label

Control actions are not supported for non-standard matches.

5.12.2.1.3Supported Rule Table Instructions

The intercepted packet is processed according to the instructions on the rule tables. The supported instructions are as follows:

- DROP - drops packet
- OFPIT_GOTO_TABLE – sends the packet for processing by another rule table

- OFPIT_METER - policer function; drops packet if it exceeds kbps/pktps limit
- OFPIT_WRITE_METADATA – writes meta-data with mask <METADATA>/0xFFF
- OFPIT_EXPERIMENTE – sends the packet for processing by another controller
- OFPIT_APPLY_ACTIONS – applies certain actions specified in the section below

5.12.2.1.4Supported ACL Apply Actions

The following actions are applied on ingress packets once a match is achieved on the ACL table:

- OFPAT_OUTPUT – the packet is sent out to a port (may also be a controller port)
- OFPAT_GROUP – the packet is sent out to a group

  3 types of group ports are supported:

  - All: The packet is broadcasted on all ports which are part of the defined group
  - Selected: The packets are distributed toward the group ports according to a weight mechanism
  - Fast-Failover (FF): FF is a group of ports, one of which is defined as the primary port through which the packets are transported. In a failure scenario (defined as part of the group definition), traffic becomes transported through the most eligible backup port (from the list of backup ports). Once the failure scenario ends, traffic is routed again through the primary port
- OFPAT_POP_VLAN – strips 802.1Q (VLAN) tag from the packet
- OFPAT_PUSH_VLAN – adds 802.1Q (VLAN) tag from the packet
- OFPAT_SET_NW_TTL – modifies network TTL
- OFPAT_DEC_NW_TTL – decrements network TTL
- OFPAT_SET_FIELD – ACL set fields detailed in section below
- Normal

5.12.2.1.5Supported ACL Set Fields

The following modifications may be implemented on ingress packets:

- OXM_OF_ETH_SRC – sets the source MAC address of the packet
- OXM_OF_ETH_DST – sets the destination MAC address of the packet
- OXM_OF_VLAN_VID – sets the VLAN ID of the packet
- OXM_OF_VLAN_PCP – sets the VLAN priority code point (PCP; 0-7)
- OXM_OF_IP_DSCP – sets IP DSCP
- OXM_OF_IP_ECN – sets network ECN

5.12.2.1.6Supported ACL Meters

- ACL tables support up to 968 meters with 1 band (drop) per meter.
- Valid meter ID range: 1-969
- Only the rate or the burst size fields can be modified using OFPMC_MODIFY

- OFPMF_BURST meter type can be OFPMF_KBPS (KB/s) or OFPMF_PKTPS (number of packets per second) but not both

  Meter actions:

- OFPMBT_DROP – drops packet according to meter configuration

5.12.2.2 FDB Table (250)

The FDB table is the same one shared with regular Onyx configuration (e.g. learning, static macs, etc). The cumulative number of supported FDB rules is 88KB. FDB may only configure rules with priority of 0x8000. Hard timeout is supported for FDB table rules. FDB rules cannot have wildcard on VID/ETH_DST.

The default action for the FDB table is normal and this cannot be changed by the user.

5.12.2.2.1 Supported FDB Apply Actions

- OFPAT_OUTPUT – the packet is sent out to a port (may be controller port)
- DROP – drops packet
- Normal

5.12.2.2.2 Supported FDB Matching Rules

- OXM_OF_VLAN_VID – matches VLAN ID
- OXM_OF_ETH_DST – matches destination MAC address

5.12.2.3 Router Table (251)

The OpenFlow router table and the regular Onyx router table share the same HW resources, but are separated logically.

The cumulative number of supported FDB & router rules is 88K. Hard timeout, where the switch removes a rule after a configured timer expires, is supported for router table rules. Switch systems ignore rule priority and configure rules according to masklen in DST IPv4/IPv6 match. A rule with action output must have SET_FIELD with ETH_DST and DEC_NW_TTL. The default action for the router table is DROP.

Set DMAC can be assigned only to one output port. When a new rule with a set DMAC and a new output port is configured, the previous rules are removed from the HW. Later, if the new configuration is deleted, the previous rules get reinstalled in HW.

Note that all sent packets from the Router Table are without a VLAN header (untagged).

5.12.2.3.1 Supported Router Apply Actions

- OFPAT_OUTPUT – the packet is sent out to a port (may be controller port)
- OFPAT_DEC_NW_TTL – decrements network TTL
- OFPAT_SET_DMAC – OFPAT_SET_FIELD with OFPXMT_OFB_ETH_DST
- DROP – drops packet

When an output action is implemented, DEC_TTL and SET_DMAC must also be set.

5.12.2.3.2Supported Router Set Fields

- OXM_OF_ETH_DST – sets the destination MAC address of the packet

5.12.2.3.3Supported Router Matching Rules

- OXM_OF_IPV4_DST – matches destination IPv4 address
- OXM_OF_IPV6_DST – matches destination IPv6 address

5.12.3 Configuring OpenFlow

➢ ***To run OpenFlow on a switch:***

Step 1. Unlock the OpenFlow CLI commands. Run:

```
switch (config) # protocol openflow
```

Step 2. Configure interfaces to be managed by OpenFlow. Run:

```
switch (config) #  1/1-1/4 openflow mode hybrid
```

Step 3. Configure the OpenFlow controller IP and TCP port. Run:

```
switch (config) # openflow controller-ip 10.209.0.205 tcp-port 6633
```

Spectrum™ based systems do not support a controller port other than the default (6633).

Step 4. (Optional) Verify the OpenFlow configuration. Run:

```
switch (config) # show openflow
OpenFlow version: OF VERSION 1.0
Table size: 1000, 0 in use
Active controller ip: 10.209.0.205 port: 6633
Connection status: HANDSHAKE_COMPLETE (CONNECTED)
Forward-to-controller: ospf lldp arp-unicast arp-broadcast (all)
Enabled ports:  Eth1/1      Eth1/2      Eth1/3      Eth1/4
switch (config) #
```

To be able to configure the switch using the controller, you should see the following line in the output:

Connection status must be: HANDSHAKE_COMPLETE (CONNECTED).

5.12.4 Configuring Flows Using CLI Commands

The on-switch commands use the Open vSwitch (OVS) syntax for OpenFlow. They are actually based on the "ovs-ofctl" command. For more details please refer to the Flow Syntax section of the following man-page: http://manpages.ubuntu.com/manpages/xenial/man8/ovs-ofctl.8.html.

It is slightly modified as you need to explicitly input a flow reference number to modify. This flow ID may be used when performing any modification to the flow (e.g. delete).

All flow configurations also appear in the running-config and are restored after switch reload.

When configuring flows, you may assign them a high priority, and then to configure a "drop all" rule for non-matching packets with a lower priority.

For the flows (use a higher priority e.g. 10000 then the drop all rule) and input interface:

```
switch (config) # openflow add-flows 1 ip, priority=5000, in_port=Eth1/1,
nw_src=192.168.0.1/32, nw_dst=239.0.1.2/32, actions=output=Eth1/56
```

The above rule matches on SRC IP=192.168.0.1 and DEST IP=239.0.1.2 and the action is to output matching traffic to interface Eth1/56.

For the "drop all" rule (use a lower priority than other match rules):

```
switch (config) # openflow add-flows 1000 priority=50,in_port=ANY,actions=DROP
```

To delete a flow, run the command "del-flows" along with a flow's reference number:

```
switch (config) # openflow del-flows 1
switch (config) # openflow del-flows 1000
```

OpenFlow may be configured using one method at a time, so if an OpenFlow controller is configured then switch CLI method cannot be used.

5.12.5 Configuring Secure Connection to OpenFlow

Since OpenFlow requires a certificate signed by the certificate authority (CA), the default certificate, which is self-signed, must be replaced.

➢ ***Changing default certificate for secure OpenFlow connection:***

Step 1. Import the certificate to be used. Run:

```
switch (config) # crypto certificate name my-openflow public-cert pem "-----BEGIN CER-
TIFICATE-----
> MIIDYzCCAksCCQC9EPbMuxjNBzANBgkqhkiG9w0BAQsFADBeMQswCQYDVQQGEwJJ
...
> fEt2ui9taB1dl9480xDsGUxwUDX4YOs/bQDjp99z+cKXUe2eYzeEwnTdrCzPZuQo
> -----END CERTIFICATE-----"
Successfully installed certificate with name 'my-openflow'
```

Step 2. Import key of certificate. Run:

```
switch (config) # crypto certificate name my-openflow private-key pem "-----BEGIN RSA
PRIVATE KEY-----
> MIIEpAIBAAKCAQEAypJnZkwbhmt71Kf/MO6cy7QmWWHhCozzWRwuWGKse+MxSmfC
...
> QAuPOVR1lSyIEnYU+X0rMHc/9tgUh/8C7mBKwj7dccMmnRWz2djsjg==
> -----END RSA PRIVATE KEY-----"
```

Step 3. Designate “my-openflow” as the global default certificate for authentication of this system to clients. Run:

```
switch (config) # crypto certificate default-cert name my-openflow
```

Step 4. Import the CA certificate which signed for the controller. Run:

```
switch (config) # # crypto certificate name rootCA public-cert pem "-----BEGIN CERTIF-
ICATE-----
> MIIDjzCCAnegAwIBAgIJALVou4mcQtxlMA0GCSqGSIb3DQEBCwUAMF4xCzAJBgNV
...
> +ZfQIOCFS8gY4BDq73W4ugr38mqIA8UXXAMPwgjCbk4NyOh0rJ1P6WT8fYzvunct
> -----END CERTIFICATE-----"
Successfully installed certificate with name 'rootCA'
```

Step 5. Adds the “rootCA” to the default CA certificate list. Run:

```
switch (config) # crypto certificate ca-list default-ca-list name rootCA
```

Step 6. Save configuration. Run:

```
switch (config) # configuration write
```

Step 7. Reboot the switch. Run:

```
switch (config) # reload
```

Step 8. Verify configuration. Run:

```
switch (config) # show crypto certificate
Certificate with name 'system-self-signed'
    Comment:                      system-generated self-signed certificate
    Private Key:                  present
    Serial Number:                0x543e2efc3a5ecdbe18b5b5e744598424
    SHA-1 Fingerprint:            14e1d36035c7a5fea9f7f0f423572c9954cb9fac

    Validity:
        Starts:                   2016/09/12 12:44:10
        Expires:                  2017/09/12 12:44:10
 Subject:
        Common Name:              switch
        Country:                  IS
        State or Province:        TBD
        Locality:                 TBD
        Organization:             TBD
        Organizational Unit:      TBD
        E-mail Address:           TBD
```

```
    Issuer:
        Common Name:              switch
        Country:                  IS
        State or Province:        TBD
        Locality:                 TBD
        Organization:             TBD
        Organizational Unit:      TBD
        E-mail Address:           TBD

Certificate with name 'my-openflow' (default-cert)
    Private Key:                  present
    Serial Number:                0xbd10f6ccbb18cd07
    SHA-1 Fingerprint:            1e0e3302182ab56f2cbd3ca21722dec55299d670

    Validity:
        Starts:                   2016/09/12 15:16:48
        Expires:                  2018/01/25 14:16:48

    Subject:
        Common Name:              switch
        Country:                  *
        State or Province:        Some-State
        Locality:                 *
        Organization:             Mlnx
        Organizational Unit:      e2e
        E-mail Address:           none@nowhere.com

    Issuer:
        Common Name:              ca
        Country:                  *
        State or Province:        Some-State
        Locality:                 *
        Organization:             Mlnx
        Organizational Unit:      e2e
    Certificate with name 'rootCA'
    Private Key:                  not present
    Serial Number:                0xb568bb899c42dc65
    SHA-1 Fingerprint:            9855536f6ee0177356ffbdc54ffe803bc83fb4c6
    Validity:
        Starts:                   2016/09/08 10:34:23
        Expires:                  2019/06/29 10:34:23

    Subject:
        Common Name:              ca
        Country:                  *
        State or Province:        Some-State
        Locality:                 *
        Organization:             Mlnx
        Organizational Unit:      e2e
```

```
    Issuer:
        Common Name:              ca
        Country:                  *
        State or Province:        Some-State
        Locality:                 *
        Organization:             Mlnx
        Organizational Unit:      e2e
```

Step 9. Configure secure controller IP connection. Run:

```
switch (config) # controller-ip 10.10.10.10 tls
```

5.12.6 Commands

protocol openflow

protocol openflow
no protocol openflow

Unhides the OpenFlow commands.
The no form of the command hides the OpenFlow commands.

Syntax Description	N/A
Default	no protocol openflow
Configuration Mode	config
History	3.3.4200
Role	admin
Example	`switch (config) # protocol openflow`
Related Commands	
Note	

openflow mode hybrid

	openflow mode hybrid **no openflow mode** Enables OpenFlow on the port. The no form of the command returns the port to its default state.
Syntax Description	N/A
Default	no openflow mode
Configuration Mode	config interface ethernet
History	3.3.4200 3.6.2100 Updated Note section
Role	admin
Example	`switch (config  1/1)# openflow mode hybrid`
Related Commands	
Note	It is possible to run “interface port-channel <port number> openflow mode hybrid”.

openflow add-flows

openflow add-flows <flow-id> [[table-id],[priority-id],<match-parameter1>[,…,< match-parameterN>],<action1>[,…,<actionN>]]

Adds OpenFlow flow.

Syntax Description	flow-id	ID number to give this flow. Range: 0-65535.
	priority-id	Priority to give this flow. Range: 0-65535.
	match-parameter	Rule according to which a match is made. For a list of supported matches, see the match column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
	table-id	Range: • ACLs: 0-249 • FDB: 250 • Router: 251 Default is 0
	action	Action to perform on the matched traffic. For a list of supported actions, see the action column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	

Example

```
switch (config  1/1)# openflow add-flows 1, priority=10,in_port=Eth1/1,
nw_src=192.168.0.1/32,nw_dst=239.0.1.2/32,actions=output=Eth 1/11,Eth 1/22,Eth 1/33

switch (config  1/1)# openflow add-flows 3 table=3,in_port=121,actions=output:117

switch (config  1/1)# openflow add-flows 2
in_port=ANY,actions=push_vlan:33024,mod_vlan_vid:4111

switch (config  1/1)# openflow add-flows 4 table=0,priority=101,
dl_type=0x0800,in_port=79,dl_vlan=233,nw_dst=172.0.0.0/8,actions=pop_vlan,goto_table:251

switch (config  1/1)# openflow add-flows 5 in_port=1,actions=dec_ttl

switch (config  1/1)# openflow add-flows 6 table=0,priority=777,
in_port=121,dl_type=0x0800,nw_proto=6,actions=mod_nw_ttl:55,output:99

switch (config  1/1)# openflow add-flows 7 table=0,priority=777,
in_port=121,dl_type=0x0800,nw_proto=6,actions=Set_field:55-\>nw_ttl,output:99
```

```
switch (config  1/1)# openflow add-flows 8 table=0,priority=777,
in_port=121,actions=output:99,Set_field:11:22:33:44:00:00-\>eth_dst

switch (config  1/1)# openflow add-flows 9 table=0,priority=777,
in_port=121,dl_type=0x0800,nw_proto=6,actions=Set_field:0-\>ip_ecn,output:99

switch (config  1/1)# openflow add-flows 10 table=0,priority=777,
in_port=121,actions=output:99,Set_field:ff:ff:ff:ff:55:66-\>eth_src

switch (config  1/1)# openflow add-flows 11 table=0,priority=777,
in_port=127,actions=group:11

switch (config  1/1)# openflow add-flows 12 priority=12,in_port=105,
actions=group:5

switch (config  1/1)# openflow add-flows 13 table=0,priority=777,
in_port=127,actions=meter:6,output:117

switch (config  1/1)# openflow add-flows 14 table=2,priority=777,
in_port=127,actions=meter:2,output:117

switch (config  1/1)# openflow add-flows 10 ip,priority=10,in_port=Eth1/1,
dl_vlan=10,actions=output=Eth1/11

switch (config  1/1)# openflow add-flows 40 ip,priority=10,in_port=Eth1/1,
action=set_field:00:0c:e9:00:00:01->eth_src,output=Eth1/11

switch (config  1/1)# openflow add-flows 30 ip,priority=100,
actions=output=normal

switch (config  1/1)# openflow add-flows 10 priority=10,in_port=ANY,
actions=DROP
```

Related Commands	
Note	If no flow-text is provided the command deletes the configured OpenFlow flows

openflow del-flows

openflow del-flows [<flow-id>]

Deletes OpenFlow flow.

Syntax Description	flow-id	ID number to give this flow. Range: 0-65535.
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow del-flows 1`	
Related Commands		
Note	If flow ID is not provided, the command deletes all configured OpenFlow flows.	

openflow add-group

openflow add-group <group-id> <group-type> <bucket-parameter1>[,…, <bucket-parameterN>]

Adds an OpenFlow group.

Syntax Description	group-id	Group ID number
	group-type	For a list of supported group types, see the group column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
	bucket parameter	Possible values: • actions=output,…,output • bucket_id=<id-number> • watch_group=<group_id> • watch_port=<port> • weight=<value>
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow add-group group_id=3, type=ff,bucket=watch_port:117,output:123,bucket=watch_port:123,output:119,bucket=watch_port:111,output:119,113,121,115,123,109,117`	
Related Commands		
Note		

openflow del-group

openflow del-group <group-id>

Deletes matching OpenFlow group ID.

Syntax Description	group-id	Group ID number
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow del-group`	
Related Commands		
Note		

openflow mod-group

openflow mod-group <group-id> <group-type> <bucket-parameter1>[,…, <bucket-parameterN>]

Modifies matching OpenFlow group ID.

Syntax Description	group-id	Group ID number
	group-type	For a list of supported group types, see the group column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
	bucket parameter	Possible values: • actions=output,…,output • bucket_id=<id-number> • watch_group=<group_id> • watch_port=<port> • weight=<value>
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow mod-group group_id=3, type=ff,bucket=watch_port:117,output:123,bucket=watch_port:123,output:119,bucket=watch_port:111,output:119,113,121,115,123,109,117,119`	
Related Commands	openflow add-group	
Note	A group must exist in order to execute this command	

openflow add-meter

openflow add-meter <meter-id> <meter-rule> <band-parameter1>[,…, <band-parameterN>]

Adds OpenFlow meter.

Syntax Description	meter-id	Meter ID number
	meter-rule	For a list of supported meters types, see the meter column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
	band-parameter	Possible values: • type={type | drop} • rate=<value> • burst_size=<size>
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow add-meter meter=6, pktps,band=type=drop,rate=10`	
Related Commands		
Note		

openflow del-meter

openflow del-meter <meter-id>

Deletes matching OpenFlow meter ID.

Syntax Description	meter-id	Meter ID number
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow del-meter meter=6`	
Related Commands		
Note		

openflow mod-meter

openflow mod-meter <meter-id> <meter-rule> <band-parameter1>[,…, <band-parameterN>]

Modifies matching OpenFlow meter ID.

Syntax Description	meter-id	Meter ID number
	meter-rule	For a list of supported meters types, see the meter column in Table 54, "OpenFlow 1.3 Pipeline Capabilities Summary Table," on page 956.
	band-parameter	Possible values: • type={type | drop} • rate=<value> • burst_size=<size>
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow mod-meter meter=6, pktps,band=type=drop,rate=10`	
Related Commands		
Note		

openflow re-apply flows

openflow re-apply flows <flow-id>

Reapplies matching flow ID.

Syntax Description	flow-id	Range: 0-65535
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow re-apply flows 58`	
Related Commands		
Note		

openflow re-apply groups

openflow re-apply groups <group-id>

Reapplies matching group ID.

Syntax Description	group-id	Range: 0-65535
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow re-apply groups group_id=2`	
Related Commands		
Note		

openflow re-apply meters

openflow re-apply meters <meter-id>

Reapplies matching meters ID.

Syntax Description	meter-id	Range: 0-65535
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Role	admin	
Example	`switch (config  1/1)# openflow re-apply meters 13`	
Related Commands		
Note		

controller-ip

openflow controller-ip <ip-address> [tls] [tcp-port <tcp-port>]
no openflow controller-ip <ip-address>

Configures the OpenFlow controller's IP & TCP port.
The command "no openflow controller-ip <ip-address>" deletes all OpenFlow controller configurations related to its IP address.
The command "no openflow controller-ip <ip-address> tcp-port" deletes all the OpenFlow controller configurations related to IP address, and any tcp-port except for TLS ones.
The command "no openflow controller-ip <ip-address> [tls] tcp-port <tcp-port>" deletes the entry for the OpenFlow controller IP address, TLS (if applicable), and the TCP port

Syntax Description	ip-address	The IPv4 address of the OpenFlow controller
	tls	Configures secure connection to OpenFlow controller
	tcp-port	Sets the TCP port number of the OpenFlow controller
Default	TCP port 6633	
Configuration Mode	config openflow	
History	3.6.1002	
	3.6.2002	Added "tls" parameter
Role	admin	
Example	`switch (config openflow) # controller-ip 10.10.10.10 tls tcp-port 6633`	
Related Commands		
Note		

datapath-id

datapath-id <value>
no datapath-id

Sets a specific identifier for the switch with which the controller is communicating. The no form of the command resets the parameter to its default value.

Syntax Description	value	The most significant 16 bits of the agent data-path ID. Range is 0x0000-0xFFFF in hexa.
Default	0x0000	
Configuration Mode	config openflow	
History	3.3.4200	
Role	admin	
Example	`switch (config openflow) # datapath-id 0x1234` `switch (config openflow) #`	
Related Commands		
Note		

openflow table match-keys

openflow table <table_id[-table_id]> match-keys <key_list>
no openflow table <table_id[-table_id]> match-keys [<key_list>]

Adds ACL keys to an OpenFlow table.
The no form of the command removes ACL keys from the OpenFlow table.

Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Valid values: 0-249.
	key_list	Key value(s)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # openflow table 1 match-keys metadata ip_proto`	
Related Commands		
Note	• OpenFlow match rules are installed according to the configured match keys • New match keys are configured only when the table is empty (i.e. does not contain any rules)	

show openflow

	show openflow Displays general information about the OpenFlow protocol configuration.
Syntax Description	N/A
Default	None
Configuration Mode	Any command mode
History	3.3.4200 3.3.4302 Removed flow-id parameter 3.6.1002 Updated Example
Role	admin

Example	`switch (config) # show openflow`
	`OpenFlow Version: OpenFlow 1.3`
	`Datapath ID: ffff7cfe90e600c0`
	`Controllers Information:`
	`Controller State Role Changed (sec) Last Error`
	`---------- ----- ---- ------------ ----------`
	`tcp:1.1.1.1:6633 BACKOFF other 3 Connection timed out`
	`tcp:10.10.10.10:6633 ACTIVE other 2067 N/A`
	`tcp:10.10.10.30:6633 ACTIVE other 2067 N/A`
	
	`Mapping of OpenFlow ports to their OpenFlow numbers:`
	`Interface OF-Port`
	`--------- -------`
	`Eth1/12 OF107`
	`Eth1/9 OF109`
	`Eth1/10 OF111`
	`Eth1/7 OF113`
	`Eth1/8 OF115`
	`Eth1/3 OF121`
	`Eth1/4 OF123`

Related Commands	
Note	

show openflow flows

	show openflow flows Displays information about the OpenFlow flows.
Syntax Description	N/A
Default	None
Configuration Mode	Any command mode
History	3.3.4302 3.6.1002 Updated Example
Role	admin

Example	`switch (config) # show openflow flows`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x0, duration=467.993s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,in_port=125 actions=output:123`
	`cookie=0x0, duration=439.218s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=9999,in_port=125 actions=output:123`
	`cookie=0x0, duration=467.984s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=1000 actions=drop`
	`cookie=0x0, duration=467.975s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=200,dl_vlan=222 actions=pop_vlan,output:123`
	`cookie=0x0, duration=467.987s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=10,dl_vlan=10 actions=output:123`
	`cookie=0x0, duration=468.013s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,dl_dst=01:01:01:01:01:01 actions=output:123`
	`cookie=0x0, duration=467.991s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,dl_src=01:01:01:01:01:01 actions=output:123`
	`cookie=0x0, duration=467.992s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=5,arp actions=output:123`

Related Commands	
Note	

show openflow flows ethernet-names

show openflow flows <cookie | table> ethernet-names

Displays OpenFlow flows configuration with interface names.

Syntax Description	N/A
Default	None
Configuration Mode	Config
History	3.6.4006
Role	admin

Example	`switch (config) # show openflow flows ethernet-names`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x0, duration=911.531s, table=0, n_packets=0, n_bytes=0,`
	`priority=0 actions=NORMAL`
	`cookie=0x0, duration=80.662s, table=1, n_packets=0, n_bytes=0,`
	`priority=0,in_port=0,dl_src=02:00:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=80.530s, table=1, n_packets=0, n_bytes=0,`
	`priority=1,in_port=1,dl_src=02:01:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=80.414s, table=1, n_packets=0, n_bytes=0,`
	`priority=2,in_port=2,dl_src=02:02:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=80.296s, table=1, n_packets=0, n_bytes=0,`
	`priority=3,in_port=3,dl_src=02:03:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=80.180s, table=1, n_packets=0, n_bytes=0,`
	`priority=4,in_port=4,dl_src=02:04:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=80.064s, table=1, n_packets=0, n_bytes=0,`
	`priority=5,in_port=5,dl_src=02:05:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=79.948s, table=1, n_packets=0, n_bytes=0,`
	`priority=6,in_port=6,dl_src=02:06:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=79.831s, table=1, n_packets=0, n_bytes=0,`
	`priority=7,in_port=7,dl_src=02:07:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=79.711s, table=1, n_packets=0, n_bytes=0,`
	`priority=8,in_port=8,dl_src=02:08:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=79.591s, table=1, n_packets=0, n_bytes=0,`
	`priority=9,in_port=9,dl_src=02:09:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`
	`cookie=0x0, duration=79.467s, table=1, n_packets=0, n_bytes=0,`
	`priority=10,in_port=10,dl_src=02:0a:00:00:00:00 actions=output:Eth1/`
	`13,output:123,output:127`

Related Commands	

Note

show openflow groups

	show openflow groups Displays all the configured OpenFlow groups.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Role	admin

Example	`switch (config) # show openflow groups`
	`OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):`
	`group_id=5566,type=select,bucket=weight:5,actions=out-put:1,bucket=weight:7,actions=output:2,bucket=weight:22,actions=out-put:3`

Related Commands	
Note	

show openflow groups ethernet-names

	show openflow groups ethernet-names Displays all the configured OpenFlow groups with their interface names.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4006
Role	admin

Example	`switch (config) # show openflow groups`
	`OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):`
	`group_id=4,type=all,bucket=actions=output:Eth1/13,output:123`
	`group_id=1,type=select,bucket=actions=output:Eth1/7,output:Eth1/8,out-put:Eth1/5,output:123,set_field:11:22:33:44:00:00->eth_dst`
	`group_id=2,type=select,bucket=actions=output:Eth1/13`
	`group_id=3,type=all,bucket=actions=output:Eth1/13,output:123,set_-field:11:22:33:44:00:00->eth_dst`

Related Commands	
Note	

show openflow meters

show openflow meters [<ID>]

Displays all/specified OpenFlow meters.

Syntax Description	ID Requested meter ID
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Role	admin

Example	`switch (config) # show openflow meters`
	`OFPST_METER_CONFIG reply (OF1.3) (xid=0x2):`
	`meter=20 kbps bands=`
	`type=drop rate=300`
	
	`meter=100 kbps bands=`
	`type=drop rate=500`
	
	`meter=200 kbps bands=`
	`type=drop rate=500`
	
	`switch (config) # show openflow meters 20`
	`OFPST_METER_CONFIG reply (OF1.3) (xid=0x2):`
	`meter=20 kbps bands=`
	`type=drop rate=300`

Related Commands	
Note	

show openflow flows table

show openflow flows table <NUM> [summary]

Displays information/summary of a given OpenFlow flows table.

Syntax Description	NUM	NUM range: 0-252
	summary	Displays given OpenFlow flow table summary
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show openflow flows table 1`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x0, duration=6.344s, table=1, n_packets=0, n_bytes=0,`
	`in_port=127 actions=drop`
	
	`switch (config) # show openflow flows table 1 summary`
	`OFPST_AGGREGATE reply (OF1.3) (xid=0x2): packet_count=0 byte_count=0`
	`flow_count=1`	

Related Commands		
Note		

show openflow flows cookie

show openflow flows cookie <cookie> [summary]

Displays information/summary of a given OpenFlow flows cookie.

Syntax Description	cookie	Requested cookie ID in the following format: cookie_id.cookie_id/mask_id (e.g. 0x2A, 0x12/0x2)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show openflow flows cookie 0x11`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x11, duration=2.699s, table=0, n_packets=0, n_bytes=0, actions=NORMAL`
	`switch (config) # show openflow flows cookie 0x22`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x22, duration=3.970s, table=1, n_packets=0, n_bytes=0, in_port=127 actions=drop`	

Related Commands		
Note	A cookie may be associated with a flow using the add-flows, and mod-flows commands.	

show openflow table match-keys

show openflow table <table_id[-table_id]> match-keys

Displays configured ACL keys in OpenFlow table.

Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Valid values: 0-249.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example		

```
switch (config) # show openflow table 2 match-keys

Table:                                                          2
Pending keys:

Key name     Description
---------    ------------
in_port      Source port
dl_src       Source MAC address
dl_dst       Destination MAC address
dl_type      Ethernet protocol type
vlan_vid     Virtual LAN tag
vlan_pcp     Priority Code Point
ip_src       Source IPv4 address
ip_dst       Destination IPv4 address
ip_proto     IPV4 - Next protocol, IPV6 - Next header
ip_dscp      IP ToS/DSCP or IPv6 traffic class field dscp
ip_ecn       ECN bits from IP header
ip_ttl       IP TTL or IPv6 hop limit
l4_src_port  Source L4 port
l4_dst_port  Destination L4 port
metadata     Matches value in the metadata field
```

Related Commands	
Note	

show openflow table match-keys supported

show openflow table <table_id[-table_id]> match-keys supported

Displays list of ACL keys which can be configured in OpenFlow table.

Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Valid values: 0-249.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example		

```
switch (config) # show openflow table 2 match-keys supported

Key name     Description
---------    ------------
in_port      Source port
dl_src       Source MAC address
dl_dst       Destination MAC address
dl_type      Ethernet protocol type
vlan_vid     Virtual LAN tag
vlan_pcp     Priority Code Point
ip_src       Source IPv4 address
ip_dst       Destination IPv4 address
ipv6_dst     Destination IPv6 address
ipv6_src     Source IPv6 address
ip_proto     IPV4 - Next protocol, IPV6 - Next header
ip_dscp      IP ToS/DSCP or IPv6 traffic class field dscp
ip_ecn       ECN bits from IP header
ip_ttl       IP TTL or IPv6 hop limit
l4_src_port  Source L4 port
l4_dst_port  Destination L4 port
metadata     Matches value in the metadata field
```

Related Commands	
Note	

5.13 VXLAN

Data centers are being increasingly consolidated and outsourced in an effort to improve the deployment time of applications and reduce operational costs, and applications are constantly raising demand for compute, storage, and network resource. Thus, in order to scale compute, storage, and network resources, physical resources are being abstracted from their logical representation, in what is referred to as server, storage, and network virtualization. Virtualization can be implemented in various layers of computer systems or networks.

Multi-tenant data centers are taking advantage of the benefits of server virtualization to provide a new kind of hosting—a virtual hosted data center. Multi-tenant data centers are ones where individual tenants could belong to a different company or a different department. To a tenant, virtual data centers are similar to their physical counterparts, consisting of end-stations attached to a network, complete with services such as load balancers and firewalls. To tenant systems, a virtual network looks like a normal network, except that the only end-stations connected to the virtual network are those belonging to a tenant's specific virtual network.

How a virtual network is implemented does not generally matter to the tenant; what matters is that the service provided (Layer 2 (L2) or Layer 3 (L3)) has the right semantics, performance, etc. It could be implemented via a pure routed network, a pure bridged network, or a combination of bridged and routed networks.

VXLAN (Virtual eXtensible Local Area Network) addresses the above requirements of the L2 and L3 data center network infrastructure in the presence of virtual networks in a multi-tenant environment. It runs over the existing networking infrastructure and provides a means to "stretch" an L2 network. Each overlay bridge is called a VXLAN segment. Only machines within the same VXLAN segment can communicate with each other. Each VXLAN segment is identified through a 24-bit segment ID called "VXLAN Network Identifier (VNI)". A network endpoint which performs a conversion from virtual to physical network and back is called VXLAN Tunnel End-Point or VTEP.

In virtual environments, it is typically required to use logical switches to forward traffic between different virtual machines (VMs) on the same physical host, between virtual machines and the physical machines and between networks. Virtual switch environments use an OVSDB management protocol for configuration and state discovery of the virtual networks. OVSDB protocol allows programmable access to the database of virtual switch configuration.

5.13.1 Configuring VXLAN

➢ ***To enable VXLAN:***

Step 1. Configure jumbo frames for NVE ports:

```
switch (config)# interface ethernet 1/1-1/4 mtu 9216 force
```

Step 2. Configure jumbo frames for underlay-facing ports:

```
switch (config)# interface ethernet 1/17 mtu 9216 force
```

Step 3. Create VLAN for all VXLAN traffic:

```
switch (config)# vlan 3
```

Step 4. Configure Overlay interfaces with VXLAN VLAN:

```
switch (config)# interface ethernet 1/17 switchport access vlan 3
```

Step 5. Enable IP routing:

```
switch (config)# ip routing vrf default
```

Step 6. Configure interface on the VXLAN VLAN and configure an IP address for it:

```
switch (config)# interface vlan 3
switch (config interface vlan 3)# ip address 33.33.33.254 255.255.255.0
switch (config interface vlan 3)# interface vlan 3 mtu 9216
```

Step 7. Enable NVE protocol:

```
switch (config)# protocol nve
```

Step 8. Configure interface NVE:

```
switch (config)# interface nve 1
```

Step 9. Create loopback interface to terminate the VXLAN tunnel. The IP address of the interface will be a VTEP endpoint address, and needs to be reachable in the underlay network:

```
switch (config)# interface loopback 1
switch (config interface loopback 1)# ip address 1.2.3.4 255.255.255.255
switch (config)# interface nve 1 vxlan source interface loopback 1
```

Step 10. Configure routing to other VTEP devices:

```
switch (config)# ip route vrf default 1.2.3.5 /32 33.33.33.253
switch (config)# ip route vrf default 1.2.3.6 /32 33.33.33.252
```

Step 11. Configure overlay-facing ports for NVE mode:

```
switch (config)#  1/1 nve mode only force
switch (config)#  1/2 nve mode only force
switch (config)#  1/3 nve mode only force
switch (config)#  1/4 nve mode only force
```

- ***For deployments with a controller, set up OVSDB:***

Step 1. Start OVSDB server:

```
switch (config)# ovs ovsdb server
```

Step 2. Configure the OVSDB manager to an IP address of a controller:

```
switch (config)# ovs ovsdb manager remote ssl ip address 10.130.250.5
```

- ***For controller-less deployments, configure the bridging from the CLI directly:***

Step 1. Create bridges:

```
switch (config)# interface nve 1 nve bridge 7777
switch (config)#  1/1 nve vlan 10 bridge 7777
```

Step 2. Configure source-node replication:

```
switch (config)# no interface nve 1 nve fdb flood load-balance
```

Step 3. Configure flood addresses for BUM traffic:

```
switch (config)# interface nve 1 nve fdb flood bridge 7777 address 1.2.3.5
switch (config)# interface nve 1 nve fdb flood bridge 7777 address 1.2.3.6
```

Step 4. Configure FDB remote learning:

```
switch (config)# interface nve 1 nve fdb learning remote
```

Figure 24: Switch Configuration & Topology

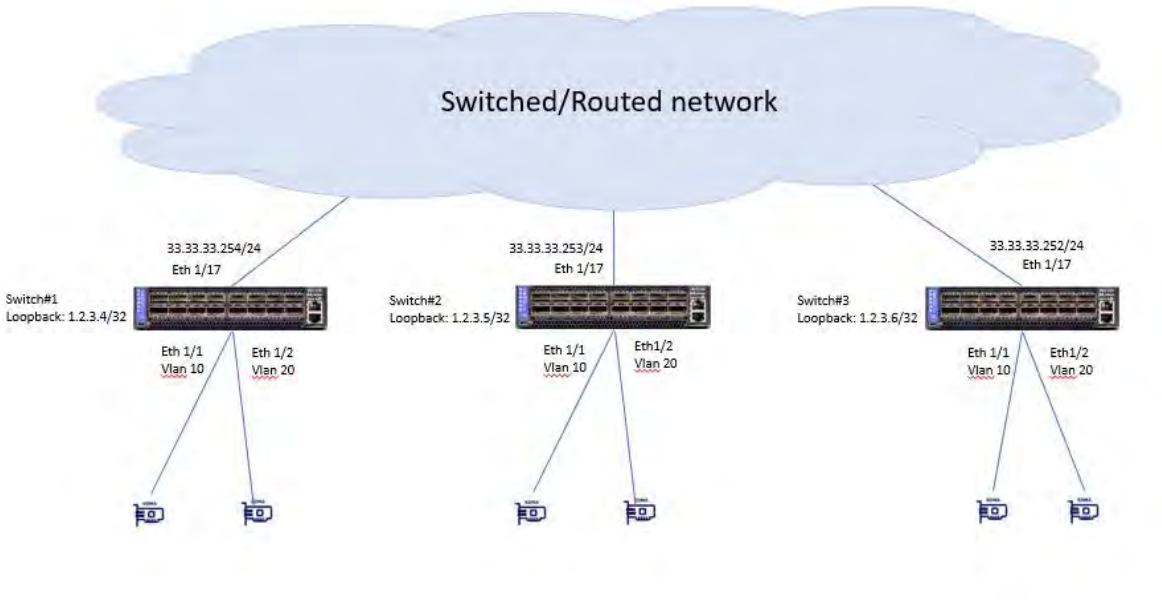

5.13.2 VMware Network Virtualization and Security Platform (NSX) Configuration

5.13.2.1 Hardware Topology

- 2 ESXi servers pre-configured with VXLAN networking using VMware NSX
- 3 NSX Controllers available for VXLAN unicast type logical switches
- 1 Mellanox switch connected to the ESXi servers and to a physical database server
- Out-of-band network for management and a VLAN network to carry VXLAN traffic

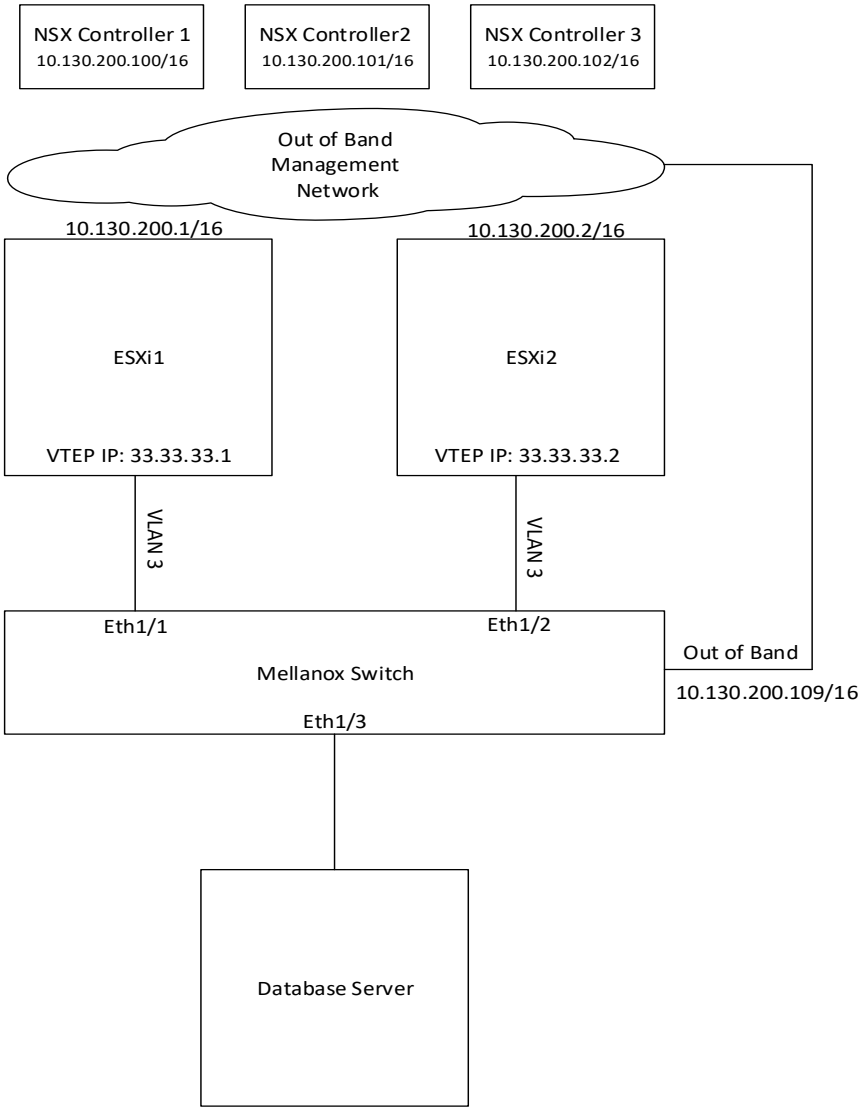

5.13.2.2 Switch Configuration

Step 1. Configure jumbo frames on ESXi and Database server facing interfaces:

```
switch (config)#  1/1-1/3 mtu 9216 force
```

Step 2. Create VLAN 3 to carry VXLAN traffic (if it does not exist yet):

```
switch (config)# vlan 3
```

Step 3. Enable IP routing:

```
switch (config)# ip routing vrf default
```

Step 4. Create an interface on VLAN 3 and assign an IP address to it.
The IP address must be the default gateway of the VXLAN netstack created by NSX after enabling VXLAN traffic on the hosts.

To check the default gateway in vSphere web client select an ESXi host and go to: Configure -> TCP/IP configuration .

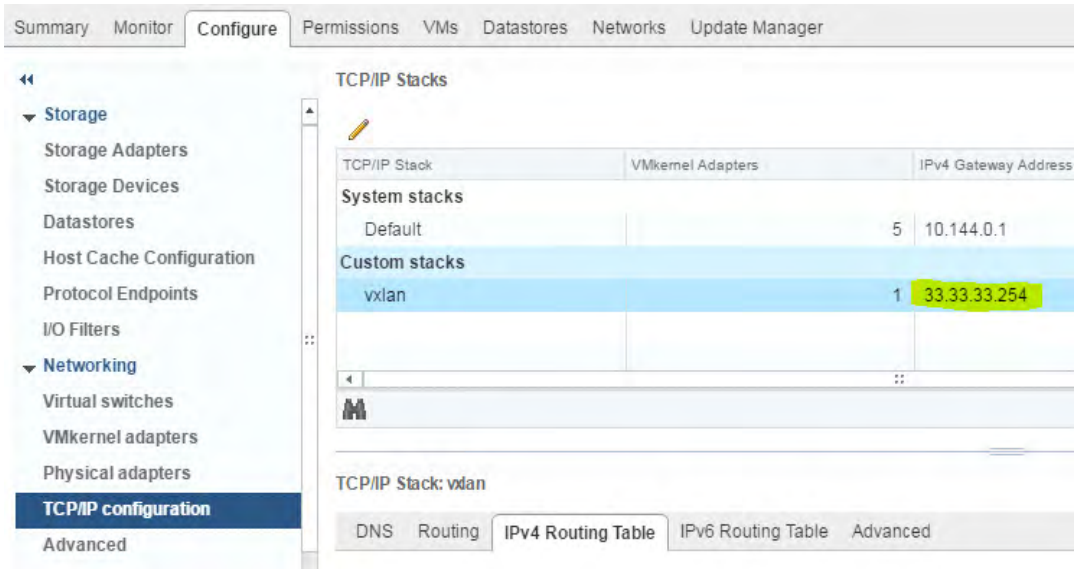


```
switch (config)# interface vlan 3
switch (config interface vlan 3)# ip address 33.33.33.254 255.255.255.0
switch (config interface vlan 3)# interface vlan 3 mtu 9216
```

Step 5. Create a loopback interface to communicate with VTEPs on the ESXi servers by routing through “interface vlan 3”. This interface will be the VTEP IP assigned to the switch

```
switch (config)# interface loopback 1
switch (config interface loopback 1)# ip address 1.2.3.4 255.255.255.255
```

Step 6. Enable NVE protocol:

```
switch (config)# protocol nve
```

Step 7. Configure interface NVE:

```
switch (config)# interface nve 1
```

Step 8. Configure the source of the NVE interface to be the loopback created above:

```
switch (config)# interface nve 1 vxlan source interface loopback 1
```

Step 9. Start the OVSDB server and connect it to the NSX Controllers

```
switch (config)# ovs ovsdb server
switch (config)# ovs ovsdb manager remote ssl ip address 10.130.200.100
switch (config)# ovs ovsdb manager remote ssl ip address 10.144.200.101
switch (config)# ovs ovsdb manager remote ssl ip address 10.144.200.102
```

Step 10. Configure the port facing the Database server as an NVE port

```
switch (config)#  1/3 nve mode only force
```

Step 11. Get the switch certificate for later configuration in the NSX Manager.

```
switch (config)# show crypto certificate name system-self-signed public-pem
```

Copy the certificate starting with the line:

```
-----BEGIN CERTIFICATE-----
```

until the line:

```
-----END CERTIFICATE-----
```

Make sure to include both of those lines.

Warning! NSX Manager Configuration

Warning! Adding Hosts to Replication Cluster

Warning! In NSX Manager, go to “Service Definitions” → “Hardware Devices”.

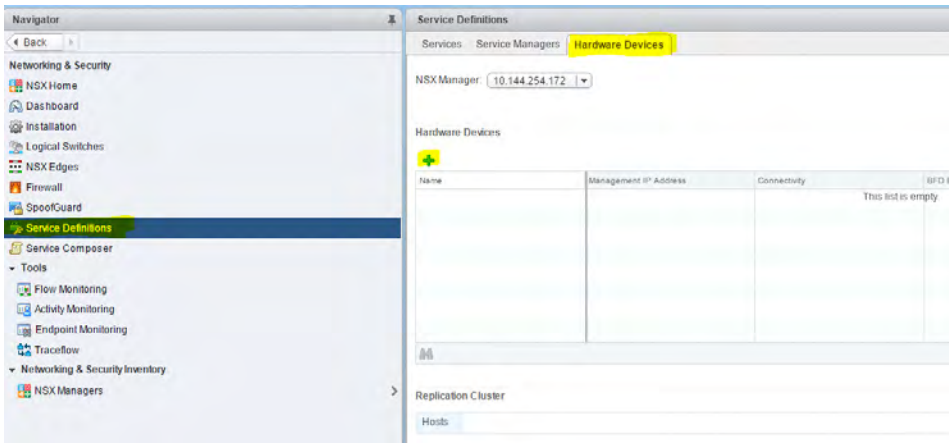

Step 12. Under “Replication Cluster” click Edit.

Step 13. Add both of the ESXi servers to the replication cluster.

All hosts added to the replication cluster can replicate BUM (Broadcast, Unknown unicast and Multicast) traffic to other ESXi servers.

When the switch needs to send BUM traffic to a virtual machine, it will select one of the hosts in the replication cluster and send the traffic to it, the host will then replicate it to all other ESXi hosts.

It is recommended to add at least 2 ESXi servers to the replication cluster for redundancy.

5.13.2.3 Adding the Mellanox Switch to NSX

Step 1. Under Hardware Devices click the + sign to add a new hardware device.

Step 2. Fill in a name for the new hardware device.

Step 3. Fill in the switch certificate we got earlier.

Step 4. Click OK.

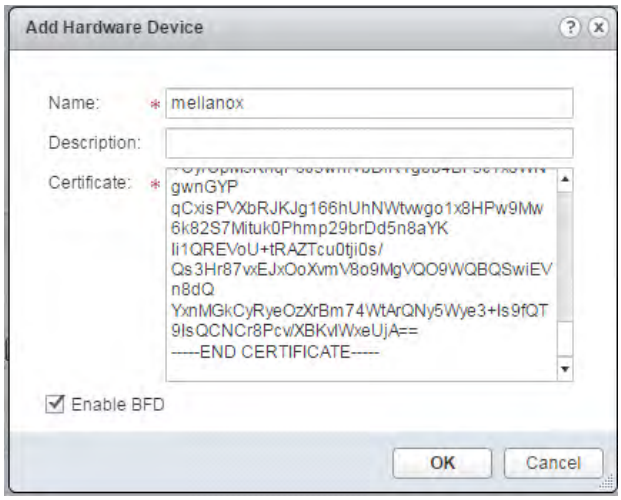

Step 5. Wait until the new switch is showing as “UP” under the connectivity column, you may need to refresh vSphere client a few times.

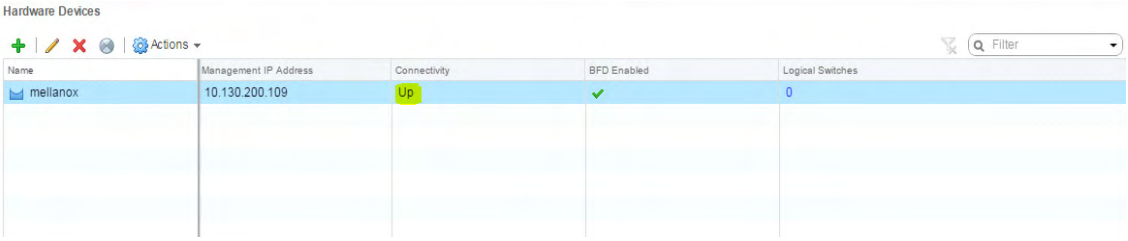

5.13.3 Mapping a Logical Switch to a Physical Switch Port

Step 1. In NSX Manager go to “Logical Switches”

Step 2. Right click the logical switch you wish to map to the physical switch port and select “Manage Hardware Bindings”

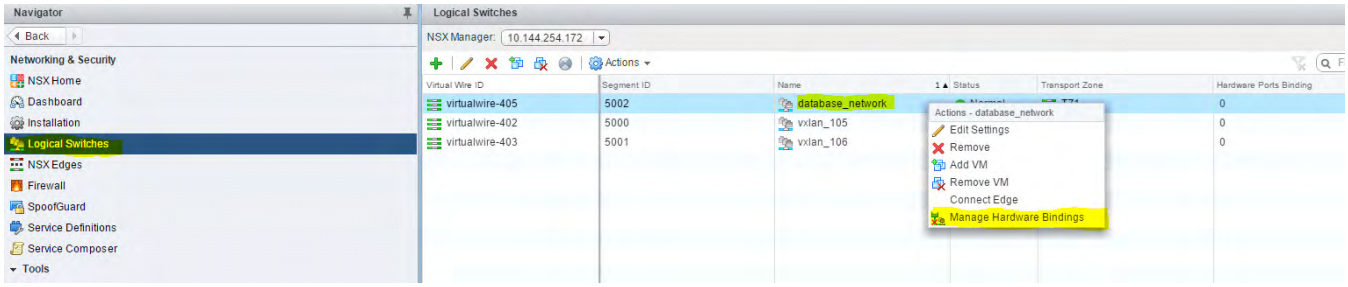

Step 3. Click the “+” sign to add a new mapping instance

Step 4. Click Select under the port column and select port “eth3”, this corresponds to “ 1/3” we configured earlier as an NVE port in the switch.

Step 5. Under the VLAN column, set the VLAN that will map this logical switch to this specific switch port, you can have multiple logical switches mapped to the same port on a different VLAN (for example to connect a firewall appliance to logical switches). For “access” con-

figuration (no VLAN is required on the host connected to the physical switch port) use VLAN 1.

Step 6. Click OK

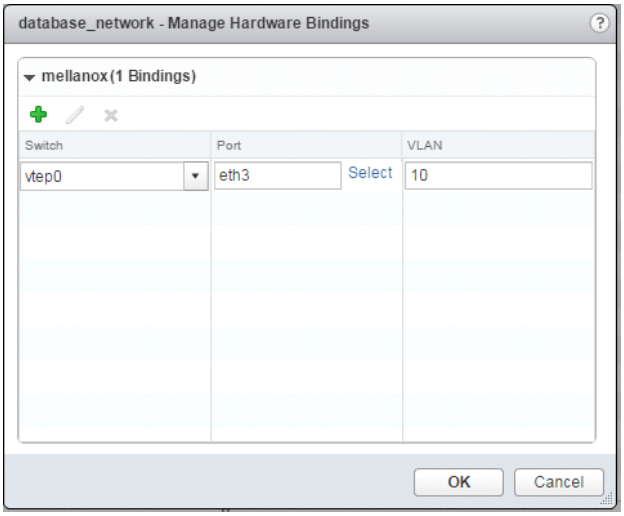

5.13.4 Commands

protocol nve

protocol nve
no protocol nve

Enables NVE functionality and displays NVE commands.
The no form of the command hides the NVE commands and deletes its database.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # protocol nve`
Related Commands	
Note	

interface nve

interface nve <nve-id>
no interface nve <nve-id>

Creates VXLAN tunnel.
The no form of the command destroys VXLAN tunnel.

Syntax Description	nve-id	NVE ID range: 1-64
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Role	admin	
Example	`switch (config) # interface nve 1` `switch (config interface nve 1) #`	
Related Commands	protocol nve	
Note		

nve bridge

nve bridge <vni-id> [name <bridge-name>]
no nve bridge <vni-id>

creates an NVE bridge with a given VNI.
The no form of the command removes NVE bridge.

Syntax Description	vni-id	VXLAN network identifier range: 0-16777216
Default	bridge-name	"bridge-<vni-id>"
Configuration Mode	config interface nve	
History	3.6.3550 3.6.3212	
Role	admin	
Example	`switch (config interface nve 1) # nve bridge 25`	
Related Commands		
Note	Number of bridges limited to 500	

nve fdb flood bridge address

nve fdb flood bridge <vni-id> address <ip-address>
no nve fdb flood bridge <vni-id> address [ip-address]

Adds an IP address of a remote VTEP to be used for BUM traffic.
The no form of the command has two input options:
• Entering an IP address removes a specific remote address
• No IP address removes all addresses

Syntax Description	vni-id	VXLAN network identifier range: 0-16777216
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.3550 3.6.3212	
Role	admin	
Example	`switch (config interface nve 1) # nve fdb flood bridge 7777 address 1.2.3.6`	
Related Commands		
Note	The number of IP addresses is limited to 750	

nve fdb flood load-balance

	nve fdb flood load-balance **no nve fdb flood load-balance** Configures service-node replication. The no form of the command configures source-node replication.
Syntax Description	N/A
Default	service-node replication
Configuration Mode	config interface nve
History	3.6.8008
Role	admin
Example	`switch (config interface nve 1) # nve fdb flood load-balance`
Related Commands	
Note	

nve fdb learning remote

	nve fdb learning remote **no nve fdb learning remote** Enables remote (controller-less) FDB learning. The no form of the command disables remote FDB learning.
Syntax Description	N/A
Default	Disabled (controller-based learning)
Configuration Mode	config interface nve
History	3.6.8008
Role	admin
Example	`switch (config interface nve 1) # nve fdb learning remote`
Related Commands	
Note	

nve mode only

nve mode only [force]
no nve mode only [force]

Sets physical interface to NVE mode.
The no form of the command removes physical interface from NVE mode.

Syntax Description	force	Forces configuration while interface is admin up
Default	Not in NVE mode	
Configuration Mode	config interface ethernet	
History	3.6.3004	
Role	admin	
Example	`switch (config  1/1) # nve mode only`	
Related Commands		
Note		

nve vlan bridge

nve vlan <vlan-id> bridge <vni-id>
no nve vlan <vlan-id> bridge <vni-id>

Maps a VLAN to a specific bridge on the interface (controller-less configuration). The no form of the command unmaps a VLAN from a specific bridge on the interface.

Syntax Description	vni-id	VXLAN network identifier range: 0-16777216
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.6102	
Role	admin	
Example	`switch (config  1/1) # nve vlan 10 bridge 7777`	
Related Commands		
Note	• Multiple VLANs cannot be mapped to a single bridge • If you use VTEP light, VLAN 1 should be used for untagged traffic	

shutdown

	shutdown **no shutdown** Disables VXLAN tunnel. The no form of the command enables VXLAN tunnel.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface nve
History	3.6.3004
Role	admin
Example	`switch (config interface nve 1) # shutdown`
Related Commands	interface nve protocol nve
Note	

vxlan source interface loopback

vxlan source interface loopback <loopback-id>
no vxlan source interface loopback <loopback-id>

Binds VXLAN tunnel to a loopback interface.
The no form of the command unbinds VXLAN tunnel from the loopback interface.

Syntax Description	loopback-id	Loopback interface ID Valid range: 0-31
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.3004	
Role	admin	
Example	`switch (config interface nve 1) # vxlan source interface loopback 14`	
Related Commands	interface nve protocol nve	
Note	The configured loopback interface becomes the VXLAN tunnel endpoint (VTEP)	

clear mac-address-table nve

clear mac-address-table nve [remote]

Clears locally-learned NVE MAC addresses.

Syntax Description	remote	Clears remotely-learned NVE MAC addresses.
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.8008	
Role	admin	
Example	`switch (config interface nve 1) # clear mac-address-table nve`	
Related Commands	interface nve protocol nve	
Note		

clear nve counters

	clear nve counters Clears NVE counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface nve
History	3.6.3004
Role	admin
Example	`switch (config interface nve 1) # clear nve counters`
Related Commands	interface nve protocol nve
Note	The command “clear counters all” also clears NVE counters

show interfaces nve

show interfaces nve [<nve-id>]

Displays information about NVE interfaces.

Syntax Description	nve-id	NVE ID range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interface nve`
	
	`Remote Manager IP Address Port Connection Type`
	`---------- ---- ---------------`
	`2.2.2.2 200 tcp`
	
	`NVE member interfaces: Eth1/2, Eth1/7`
	
	`Interface NVE 1 status:`
	`Admin state: up`
	`Source interface: loopback 1`
	
	`17971 encapsulated (Tx) NVE packets`
	`0 decapsulated (Rx) NVE packets`
	`0 dropped NVE-encapsulated packets`
	`0 NVE-encapsulated packets with errors`	

Related Commands		
Note		

show interfaces nve counters

show interfaces nve <nve-id> counters

Displays NVE counters.

Syntax Description	nve-id	NVE ID range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interface nve 1 counters`
	`  18330  encapsulated (Tx) NVE packets`
	`  0  decapsulated (Rx) NVE packets`
	`  0  dropped NVE-encapsulated packets`
	`  0  NVE-encapsulated packets with errors`	

Related Commands		
Note		

show interfaces nve flood

show interfaces nve <nve-id> flood [vni <vni-id>]

Displays remote VTEP endpoints configured for BUM (broadcast, unknown unicast, multicast) flooding.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interface nve 1 flood`
	
	`NVE Interface  Logical Switch  VNI ID  Flood IP Address`
	`-------------  --------------  ------  ----------------`
	`1  ls7777  7777  1.2.3.5`	

Related Commands		
Note		

show interfaces nve mac-address-table

show interfaces nve <nve-id> mac-address-table [vni <vni-id>]

Displays MAC address table of NVE interface.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays MAC address table of NVE interface with specified VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show interface nve 1 mac-address-table`	

```
NVE Interface   Logical Switch       VNI ID    Mac Address        Address Type       Remote Endpoint IP Address
-------------   --------------       ------    -----------        ------------       --------------------------
1               ls7777               7777      e4:1d:2d:a5:f2:0a  local learned      N/A
1               ls7777               7777      00:11:22:33:44:55  remote configured  1.2.3.5
```

Related Commands	
Note	

show interfaces nve mac-address-table local learned unicast

show interfaces nve <nve-id> mac-address-table local learned unicast [vni <vni-id>]

Displays only the locally-learned unicast MAC addresses.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays MAC addresses on the bridge with the given VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show interface nve 1 mac-address-table local learned unicast`	

```
NVE Interface   Logical Switch       VNI ID    Mac Address        Address Type       Remote Endpoint IP Address
-------------   --------------       ------    -----------        ------------       --------------------------
1               ls7777               7777      e7:3a:7e:a5:f2:1a  local learned      N/A
```

Related Commands	
Note	

show interfaces nve mac-address-table remote configured multicast

show interfaces nve <nve-id> mac-address-table remote configured multicast [vni <vni-id>]

Displays only remotely-configured BUM addresses.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays only MAC addresses on the bridge with the given VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show interface nve 1 mac-address-table remote configured multicast`	

```
NVE Interface   Logical Switch       VNI ID    Mac Address        Address Type       Remote Endpoint IP Address
-------------   --------------       ------    -----------        ------------       --------------------------
1               ls7777               7777      00:11:22:33:44:55  remote configured  1.2.3.5
```

Related Commands	
Note	

show interfaces nve mac-address-table remote configured unicast

show interfaces nve <nve-id> mac-address-table remote configured unicast [vni <vni-id>]

Displays only remotely-configured unicast addresses.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays only MAC addresses on the bridge with the given VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	
Example	`switch (config) # show interface nve 1 mac-address-table remote configured unicast`	

```
NVE Interface   Logical Switch       VNI ID    Mac Address        Address Type       Remote Endpoint IP Address
-------------   --------------       ------    -----------        ------------       --------------------------
1               ls7777               7777      00:11:22:33:44:55  remote configured  1.2.3.5
```

Related Commands	
Note	

show interfaces nve peers

show interfaces nve <nve-id> peers [vni <vni-id>]

Displays all remote VTEPs.

Syntax Description	nve-id	NVE ID range: 1-64
	vni	Displays NVE peers on specific VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interface nve 1 peers`
	
	`NVE Interface  Logical Switch  VNI ID  Peer IP Address`
	`-------------  --------------  ------  ----------------`
	`1  ls7777  7777  1.2.3.5`	

Related Commands		
Note		

ovs ovsdb server

	ovs ovsdb server **no ovs ovsdb server** Runs OVSDB-server process and unhides OVS commands. The no form of the command deactivates OVSDB-server process and hides OVS commands.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config) # ovs ovsdb server`
Related Commands	
Note	OVSDB server runs when “protocol openflow” or “protocol nve” are enabled, even when not enabled using this command

ovs ovsdb manager remote

ovs ovsdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port>
no ovs ovsdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port>

Configures OVSDB to actively connect to a remote manager at a given IP address and TCP port, using either TCP or SSL.
The no form of the command disconnects OVSDB from a remote manager.

Syntax Description	SSL	Connect with TCP protocol
	TCP	Connect with SSL protocol
	ip-address	IP address of remote manager
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Role	admin	
Example	`switch (config) # ovs ovsdb manager remote tcp ip-address 10.10.10.10 port 20`	
Related Commands	ovs ovsdb server	
Note		

ovs ovsdb server listen

ovs ovsdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address>
no ovs ovsdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address>

Configures OVSDB to listen at a given port of an interface with a given (local) IP address.
The no form of the command disconnects OVSDB from a remote manager.

Syntax Description	SSL	Connect with TCP protocol
	TCP	Connect with SSL protocol
	ip-address	IP address of a given port
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Role	admin	
Example	`switch (config) # ovs ovsdb server listen tcp port 20 local ip-address 20.20.20.20`	
Related Commands	ovs ovsdb server	
Note		

5.14 IGMP Snooping

The Internet Group Multicast Protocol (IGMP) is a communications protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships. The host joins a multicast-group by sending a join request message towards the network router, and responds to queries sent from the network router by dispatching a join report.

A given port can be either manually configured to be a MRouter port or it can be dynamically manifested when having received a query, hence, the network router is connected to this port. All IGMP Snooping control packets received from hosts (joins/leaves) are forwarded to the MRouter port, and the MRouter port updates its multicast-group data-base accordingly. Each dynamically learned multicast group will be added to all of the MRouter ports on the switch.

As many as 5K multicast groups can be created on the switch.

5.14.1 Configuring IGMP Snooping

You can configure IGMP snooping to establish multicast group memberships.

➢ ***To configure IGMP snooping:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
```

Step 4. Enable IGMP snooping on a VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping
```

5.14.2 Defining a Multicast Router Port on a VLAN

You can define a Multicast Router (MRouter) port on a VLAN in one of the following methods:

➢ ***To change the interface switchport to trunk:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
switch (config) #
```

Step 4. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) #  1/1
switch (config  1/1) # switchport mode trunk
```

Step 5. Change back to config mode. Run:

```
switch (config  1/1) # exit
switch (config) #
```

Step 6. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping mrouter  1/1
switch (config vlan 2) #
```

➢ ***To change the interface switchport to hybrid:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
switch (config) #
```

Step 4. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) #
```

Step 5. Change back to config mode. Run:

```
switch (config vlan 200) # exit
switch (config) #
```

Step 6. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) #  1/22
switch (config  1/22) # switchport mode hybrid
```

Step 7. Attach the VLAN to the port's interface. Run:

```
switch (config  1/22) # switchport mode hybrid allowed-vlan 200
switch (config  1/22) #
```

Step 8. Change to config mode again. Run:

```
switch (config  1/22) # exit
switch (config) #
```

Step 9. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # ip igmp mrouter  1/22
switch (config vlan 200) #
```

5.14.3 IGMP Snooping Querier

IGMP Snooping Querier complements the IGMP snooping functionality. IGMP Snooping Querier is used to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not need to be routed. When IGMP Snooping Querier is

enabled, IGMP queries are sent out periodically by the switch through all ports in the VLAN and to which hosts wishing to receive IP multicast traffic respond with IGMP report messages. IGMP Snooping Querier must be used in conjunction with IGMP snooping as IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

➢ ***To configure IGMP Snooping Querier:***

Step 1. Enable the IGMP snooping on the switch. Run:

```
switch (config) # ip igmp snooping
```

Step 2. Enable the IGMP snooping querier on a specific VLAN. Run:

```
switch (config) # vlan 10
switch (config vlan 10)# ip igmp snooping querier
```

Step 3. Set the query interval time. Run:

```
switch (config vlan 10)# ip igmp snooping querier query-interval 100
```

Step 4. (Optional) Verify the IGMP snooping querier configuration. Run:

```
switch (config vlan 10)# show ip igmp snooping querier
Snooping querier information for VLAN 10

IGMP Querier Present
Querier IP address: 1.1.1.2
Query interval: 125
Response interval: 100
Group membership interval: 1
Robustness: 2
Version: 2

switch (config vlan 10)#
```

5.14.4 Commands

ip igmp snooping (admin)

	ip igmp snooping **no ip igmp snooping** Enables IGMP snooping globally or per VLAN. The no form of the command disables IGMP snooping globally or per VLAN.
Syntax Description	N/A
Default	IGMP snooping is disabled, globally and per VLAN.
Configuration Mode	config config vlan
History	3.1.1400
Role	admin

Example	`switch (config) # ip igmp snooping`
	`switch (config) # vlan 10`
	`switch (config vlan 10) # ip igmp snooping`

Related Commands	show ip igmp snooping
Note	IGMP snooping has global admin state, and per VLAN admin state. Both states need to be enabled in order to enable the IGMP snooping on a specific VLAN.

ip igmp snooping (config)

ip igmp snooping {last-member-query-interval <1-25> | proxy reporting mrouter-timeout <60-600> | port-purge-timeout <130-1225> | report-suppression-interval <1-25> | unregistered multicast {flood | forward-to-mrouter-ports} | version {2 | 3}}
no ip igmp snooping {last-member-query-interval | proxy reporting | mrouter-timeout | report-suppression-interval | unregistered multicast | version}

Configures global IGMP parameters.
The no form of the command resets the global IGMP parameters to default.

Syntax Description	last-member-query-interval <1-25>	Sets the time period (in seconds) with which the general queries are sent by the IGMP querier. After timeout expiration, the port is removed from the multicast group.
	proxy reporting	Enables proxy reporting
	mrouter-timeout <60-600>	Sets the IGMP snooping MRouter port purge time-out after which the port gets deleted if no IGMP router control packets are received
	port-purge-timeout <130-1225>	Sets the IGMP snooping port purge time interval after which the port gets deleted if no IGMP reports are received
	report-suppression-interval <1-25>	Sets the IGMP snooping report-suppression time interval for which the IGMPv2 report messages for the same group will not get forwarded onto the MRouter ports
	unregistered multicast	Sets the behavior of the snooping switch for unregistered multicast traffic
	version	Sets the default operating version to use for newly created IGMP snooping instances • 2 – enables IGMPv2 • 3 – enables IGMPv3 Also available in “config vlan” configuration mode
Default	last-member-query-interval – 1 second proxy reporting is disabled mrouter-timout – 125 port-purge-timeout – 260 seconds report-suppression-interval – 5 seconds unregistered multicast – flood version – 3	
Configuration Mode	config	

History	3.1.1400	
	3.2.0500	Added “unregistered multicast” parameter
	3.6.1002	Added “version parameter”
	3.6.2100	Changed default value for “version” parameter
	3.7.11xx	Updated note
Role	admin	
Example	`switch (config) # ip igmp snooping report-suppression-interval 3`	
Related Commands	ip igmp snooping (admin) show ip igmp snooping	
Note	• When both 'IGMP and IGMP Snooping' protocols handle a Leave message and have different values for "Last Member Query Time" timer configured, then there is traffic loss for a short period of time.	

ip igmp snooping fast-leave

	ip igmp snooping fast-leave **no ip igmp snooping fast-leave** Enables fast leave processing on a specific interface. The no form of the command disables fast leave processing on a specific interface.
Syntax Description	N/A
Default	Normal-leave is enabled.
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.1.1400 3.3.4500 Added MPO configuration mode
Role	admin
Example	`switch (config  1/1) # ip igmp snooping fast-leave`
Related Commands	show ip igmp snooping interfaces
Note	

ip igmp snooping mrouter

ip igmp snooping mrouter interface <type> <number>
no ip igmp snooping mrouter interface <type> <number>

Creates a static multicast router port on a specific VLAN, on a specific interface. The no form of the command removes the static multicast router port from a specific VLAN.

Syntax Description	interface <type> <number>	Attaches the group to a specific interface. type - ethernet or port-channel.
Default	No static mrouters are configured	
Configuration Mode	config vlan	
History	3.1.1400	
Role	admin	
Example	`switch (config)# vlan 1` `switch (config vlan 1) # ip igmp snooping mrouter  1/1`	
Related Commands	show ip igmp snooping mrouter	
Note	The multicast router port can be created only if IGMP snooping is enabled both globally and on the VLAN.	

ip igmp snooping static-group

ip igmp snooping static-group <IP address> interface <type> <number> [source <source-IP>]
no ip igmp snooping static-group <IP address> interface <type> <number> [source <source-IP>]

Creates a specified static multicast group for specified ports and from a specified source IP address.
The no form of the command deletes the interface from the multicast group.

Syntax Description	IP address	Multicast IP address <224.x.x.x - 239.255.255.255>
	interface	Attach the group to a specific interface
	type	Ethernet or port-channel
	source	Source IP address If omitted, a multicast group is created for all sources
Default	No static groups are configured	
Configuration Mode	config vlan	
History	3.1.1400	
	3.6.2100	Added “source” parameter
Role	admin	
Example	`switch (config vlan 1) # ip igmp snooping static-group 230.0.0.1  1/1`	
Related Commands	show ip igmp snooping groups	
Note	If the deleted interface is the last port, it deletes the entire multicast group.	

ip igmp snooping querier

	ip igmp snooping querier **no ip igmp snooping querier** Enables the IGMP Snooping Querier on a VLAN. The no form of the command disables the IGMP Snooping Querier on a VLAN.
Syntax Description	N/A
Default	Disable
Configuration Mode	config vlan
History	3.3.4200
Role	admin
Example	`switch (config vlan 1)# ip igmp snooping querier` `switch (config vlan 1)#`
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier
Note	

igmp snooping querier query-interval

igmp snooping querier query-interval <time>
no igmp snooping querier query-interval

Configures the query interval.
The no form of the command rests the parameter to its default.

Syntax Description	time	Time interval between queries (in seconds).
Default	125 seconds	
Configuration Mode	config vlan	
History	3.3.4200	
	3.7.1000	Updated example
Role	admin	
Example	`switch (config vlan 1)# igmp snooping querier query-interval 100`	
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier	
Note		

clear ip igmp snooping counters

clear ip igmp snooping counters [vlan <vlan-id>]

Clears IGMP snooping counters.

Syntax Description	vlan	Clears IGMP snooping counters per VLAN
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.6.6000	Updated command format
Role	admin	
Example	`switch (config) # clear ip igmp snooping counters vlan 2`	
Related Commands		
Note		

show ip igmp snooping

	show ip igmp snooping Displays IGMP snooping information for all VLANs or a specific VLAN.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.6.1002 Added default IGMP version to Example 3.6.6102 Updated Example
Role	admin

Example	`switch (config) # show ip igmp snooping`
	
	`IGMP snooping global configuration:`
	`  IGMP snooping globally: enabled`
	`  IGMP default version for new VLAN: V3`
	`  IGMP snooping operationally: enabled`
	`  Proxy-reporting globally: enabled`
	`  Last member query interval: 1 seconds`
	`  Mrouter timeout: 125 seconds`
	`  Port purge timeout: 260 seconds`
	`  Report suppression interval: 5 seconds`
	`  IGMP snooping unregistered multicast: flood`

Related Commands	
Note	

show ip igmp snooping groups

show ip igmp snooping groups [vlan <vlan ID> [group <group IP >]]

Displays per VLAN the list of multicast groups attached (static or dynamic allocated) per port.

Syntax Description	vlan	VLAN ID
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.1002	Updated Example
	3.6.2100	Added “vlan” and “group” parameters and updated example
	3.6.6102	Updated example output
Role	admin	

Example

```
switch (config) # show ip igmp snooping groups
-------------------------------------------------
Vlan ID     Group           St/Dyn     Ports
-------------------------------------------------
1           230.0.0.1       St         Eth1/1,Eth1/2
2           230.0.0.1       St         Eth1/4,Eth1/6
2           230.0.0.2       St         Eth1/5

Total Num of Dynamic Group Addresses: 1
Total Num of Static Group Addresses: 1

switch (config) # show ip igmp snooping groups vlan 1
-------------------------------------
Group           St/Dyn     Ports
-------------------------------------
230.0.0.1       St         Eth1/1,Eth1/2,Eth1/3

Total Num of Dynamic Group Addresses: 0
Total Num of Static Group Addresses: 1

switch (config) # show ip igmp snooping groups vlan 1 group 230.0.0.1
Snooping group information for VLAN 1 and group 230.0.0.1
        Filter Mode: EXCLUDE
        Exclude sources: None
        V1/V2 Receiver Ports: Eth1/1,Eth1/2,Eth1/3
        V3 Receiver Ports:   None
```

Related Commands	
Note	

show ip igmp snooping interfaces

	show ip igmp snooping interfaces Displays IGMP snooping interface information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400
Role	admin

Example	`switch (config) # show ip igmp snooping interfaces`
	`interface  leave-mode`
	`-----------  ------------`
	`1/1  Normal`
	`1/2  Normal`
	`1/3  Normal`
	`1/4  Fast`
	`...`
	`switch (config) #`

Related Commands	
Note	

show ip igmp snooping membership

show ip igmp snooping membership [vlan <VID> [group <group IP>]]

Displays IGMP snooping querier counters.

Syntax Description	vlan	Displays IGMP snooping querier counters on specific VLAN
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2100	
Role	admin	

Example	`switch (config) # show ip igmp snooping membership vlan 1 group 224.5.5.5`
	`Snooping membership information for VLAN 1 and group 224.5.5.5`
	
	`Receiver Port: Eth1/1`
	`  Attached Host: 10.10.10.1`
	`    Version: 3`
	`    Mode: Include`
	`    Sources: 10.10.10.100`
	`    Timeout since the host has been joined: 0:00:02`
	`    Expiry timeout: 0:04:18`	

Related Commands		
Note		

show ip igmp snooping mrouter

	show ip igmp snooping mrouter Displays IGMP snooping multicast router information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400
Role	admin

Example	`switch (config) # show ip igmp snooping mrouter`
	`Vlan  Ports`
	`--------  ------------`
	`1  Eth1/1(static)`
	`switch (config) #`

Related Commands	
Note	

show ip igmp snooping querier

show ip igmp snooping querier [vlan <num>]

Displays running IGMP snooping querier configuration on the VLANs.

Syntax Description	vlan <num>	Displays the IGMP snooping querier configuration running on the specified VLAN.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4200	
	3.6.2100	Updated Example
Role	admin	

Example	`switch (config) # show ip igmp snooping querier vlan 1`
	`Snooping querier information for VLAN 1`
	
	`IGMP Querier Present`
	`Querier IP address: 10.10.10.10`
	`Query interval: 125`
	`Response interval: 100`
	`Group membership interval: 1`
	`Robustness: 2`
	`Version: 3`	

Related Commands		
Note		

show ip igmp snooping querier counters

show ip igmp snooping querier counters [vlan <num> [group <group-id>]]

Displays IGMP snooping querier counters.

Syntax Description	vlan	Displays IGMP snooping querier counters on specific VLAN
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Role	admin	

Example	`switch (config) # show ip igmp snooping querier counters vlan 10`
	`Snooping querier counters for VLAN 10`
	`General queries received: 0`
	`General queries transmitted: 0`
	`Group specific queries received : 0`
	`Group specific queries transmitted : 0`
	`Group source specific queries received : 0`
	`Group source specific queries transmitted : 0`
	`Leave messages received : 0`
	`Leave messages transmitted : 0`
	`V1/V2 reports received : 0`
	`V1/V2 reports transmitted : 0`
	`V3 reports received: 0`
	`V3 reports transmitted: 0`	

Related Commands		
Note		

show ip igmp snooping statistics

	show ip igmp snooping statistics Displays IGMP snooping statistical counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400 3.6.1002 Updated Example 3.6.2100 Updated Example
Role	admin

Example	`switch (config) # show ip igmp snooping statistics`
	`Snooping Statistics for VLAN 3770`
	`General queries received : 3`
	`General queries transmitted: 0`
	`Group specific queries received : 0`
	`Group specific queries transmitted: 0`
	`Group and source specific queries received : 0`
	`Group and source specific queries transmitted: 0`
	`V1/V2 reports received : 0`
	`V1/V2 reports transmitted : 0`
	`Leave messages received : 0`
	`Leave messages transmitted: 0`
	`V3 reports received : 12`
	`V3 reports transmitted : 0`
	`Active Groups count: 2`
	`Dropped packets: 0`
	`Joins: 0`

Related Commands	
Note	

show ip igmp snooping vlan

show ip igmp snooping vlan {<vlan/vlan-range> | all}

Displays IGMP configuration per VLAN or VLAN range.

Syntax Description	vlan/vlan range	Displays IGMP VLAN configuration per specific VLAN or VLAN range.
	all	Display IGMP VLAN configuration on all VLAN.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
Role	admin	

Example	`switch (config) # show ip igmp vlan 1`
	`Vlan 1 configuration parameters:`
	`  IGMP snooping is enabled`
	`  IGMP version is V2`
	`  Snooping switch is acting as Non-Querier`
	`  mrouter static port list: Eth1/1`
	`  mrouter dynamic port list: none`	

Related Commands		
Note		

5.15 Priority Flow Control

Priority Flow Control (PFC) provides an enhancement to the existing pause mechanism in Ethernet. The current Ethernet pause option stops all traffic on a link. PFC creates eight separate virtual links on the physical link and allows any of these links to be paused and restarted independently, enabling the network to create a no-drop class of service for an individual virtual link.

PFC offers the following features:

- Provides per-priority enabling or disabling of flow control
- Transmits PFC-PAUSE frames when the receive threshold for a particular traffic class is reached
- Provides the management capability for an administrator to configure the flow control properties on each port of the switch
- Keeps flow control disabled for all priorities on all ports by default
- Allows an administrator to enable or disable flow control per port and per priority level
- Supports flow control only on physical ports, not on logical interfaces such as tunnels or interfaces defined by sharing a physical port in multiple virtual switch contexts
- Uses the configured threshold values to set up the queue buffer spaces accordingly in the data-path
- Provides hardware abstraction layer call-outs for the following:
  - Enabling or disabling of flow control on each port for each priority
  - Configuring the queue depth for each priority on each port
- Provides trace logs for execution upon error conditions and for any event notifications from the hardware or data-path. These trace logs are a useful aid in troubleshooting.
- Allows the administrator to configure the minimum and maximum threshold values for flow control. These configurations are applied globally on all ports and priorities.

Priority Based Flow Control (PFC) provides an enhancement to the existing pause flow control mechanism as described in 802.1x.

➢ ***To enable PFC globally:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable PFC globally on the switch. Run:

```
switch (config) # dcb priority-flow-control enable
This action might cause traffic loss while shutting down a port with priority-flow-con-
trol mode on
Type 'yes' to confirm  enable pfc globally: yes
```

➢ ***To enable PFC per priority:***

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable PFC globally on the switch. Run:

```
switch (config) # dcb priority-flow-control enable
# dcb priority-flow-control enable
This action might cause traffic loss while shutting down a port with priority-flow-con-
trol mode on
Type 'yes' to confirm  enable pfc globally: yes
switch (config) #
```

Step 4. Choose the priority you want to enable using the command `dcb priority-flow-control priority <pri[0..7]> enable`.

```
switch (config) # dcb priority-flow-control priority 5 enable
```

➢ ***To enable PFC per interface:***

Step 1. Log in as admin.

Step 2. Change to config mode. Run:

```
switch > enable
switch # configure terminal
```

Step 3. Enable PFC globally on the switch. Run:

```
switch (config) # dcb priority-flow-control enable
```

Step 4. Choose the priority you want to enable using the command `dcb priority-flow-control priority <pri[0..7]> enable`

```
switch (config) # dcb priority-flow-control 5 enable
```

Step 5. Change to Interface mode. Run:

```
switch (config) #
switch (config) #  1/1
switch (config  1/1) #
```

Step 6. Enable PFC for the specific interface:

```
switch (config  1/1) # dcb priority-flow-control mode on
```

When working with lossless traffic, the receiving side sends a pause frame (Xoff) to the transmitting side before the buffer is filled. When the buffer empties, the receiving side sends an un-pause frame (Xon) to the transmitting side.

5.15.1 Flow Control Threshold Configuration

The user has to set the buffer usage Xoff and Xon thresholds. The thresholds depend on network parameters (bandwidth, link latency, MTU) and the allocated size for the region.

Figure 25: Xon/Xoff Configuration

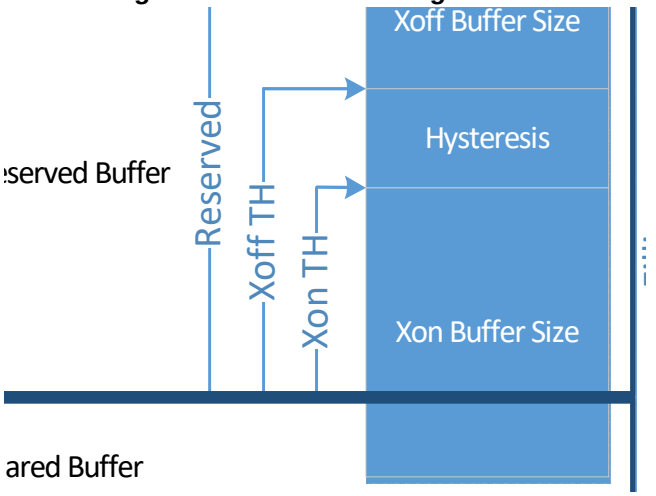

When working with global flow control mode only, a single PG shall be used and Xoff and Xon shall be set on this PG. When working with priority flow control, Xoff and Xon shall be set on each lossless PG.

See Section 5.17, “Shared Buffers,” on page 1106 for more information on flow control.

5.15.2 PFC Watchdog

Lossless networks with PFC enabled provide strong packet delivery guarantees. However, lossless networks introduce a new fault scenario where a queue of an end-port (e.g. the port of a host connected to the network) may not be able to receive any traffic from the network and keeps sending pause frames towards the switch. Since lossless switch paths do not drop packets but decline receiving more packets when their buffers fill up, if the end-port queue is stuck for a long time, the buffers fill up not only for the target switch, but also on all switches with problematic port queues in the traffic forwarding path. This leads to endless PFC pause frames, also called a PFC storm, being observed on all switch ports along the path to the traffic source.

PFC watchdog prevents congestion from spreading in such a case. When switches detect this situation on any TC queue, all the packets in the queue are flushed and new packets destined to the same queue are dropped as well until PFC storming is relieved.

5.15.3 Commands

dcb priority-flow-control enable

dcb priority-flow-control enable [force]
no dcb priority-flow-control enable [force]

Enables PFC globally on the switch.
The no form of the command globally disables PFC on the switch.

Syntax Description	force	Forces operation
Default	PFC is disabled.	
Configuration Mode	config	
History	3.1.0000	
	3.3.0000	Updated Example
Role	admin	

Example	`switch (config)# dcb priority-flow-control enable`
	`This action might cause traffic loss while shutting down a port with priority-flow-control mode on`
	`Type 'yes' to confirm enable pfc globally: yes`	

Related Commands	show dcb priority-flow-control	
Note	This command asks the user to approve traffic loss because some interfaces with DCB mode activated might get shut down.	

dcb priority-flow-control priority

dcb priority-flow-control priority <prio> enable
no dcb priority-flow-control priority <prio> enable

Enables PFC per priority on the switch.
The no form of the command disables PFC per priority on the switch.

Syntax Description	prio 0-7.
Default	PFC is disabled for all priorities.
Configuration Mode	config
History	3.1.0000
Role	admin
Example	`switch (config)# dcb priority-flow-control priority 0 enable`
Related Commands	show dcb priority-flow-control
Note	

dcb priority-flow-control mode

dcb priority-flow-control mode <mode> [force]
no dcb priority-flow-control mode [force]

Changes PFC mode per interface.
The no form of the command disables PFC per interface.

Syntax Description	force	Configures the PFC admin mode as on or auto with no confirmation needed if the port is admin enabled
	mode	The interface PFC mode. Possible values: • on – enables PFC per interface • off – disables PFC per interface • auto – set PFC mode for the interface to be controlled with traffic pool configuration
Default	auto – PFC mode is established by traffic pool configuration (not a directly configurable mode)	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
	3.6.6000	Added “force” parameter
	3.6.6102	Added “mode” parameter
	3.6.7100	Updated “mode” parameter description
Role	admin	
Example	`switch (config  1/1) # dcb priority-flow-control mode on`	
Related Commands	show dcb priority-flow-control	
Note	• For the “force” parameter, the no form of the command disables priority-flow-control without the preceding confirmation prompt • For mode value “auto”, if a lossless traffic pool is configured, PFC is enabled for this port. Otherwise, PFC is disabled	

pfc-wd

	pfc-wd **no pfc-wd** Enables PFC watchdog on interface. The no form of the command disables PFC watchdog on interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.6000
Role	admin
Example	`switch (config  1/1) # pfc-wd`
Related Commands	show interface pfc-wd
Note	

show dcb priority-flow-control

show dcb priority-flow-control [interface <type> <inf>] [detail]

Displays DCB priority flow control configuration and status.

Syntax Description	type	• ethernet • port-channel
	inf	The interface number.
	detail	Adds details information to the show output.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Role	admin	

Example

```
switch (config  1/1) # show dcb priority-flow-control

PFC enabled
Priority Enabled List    : 0
Priority Disabled List   : 1 2 3 4 5 6 7

TC     Lossless
---    ----------
0           N
1           Y
2           Y
3           N


Interface      PFC admin        PFC oper
------------   --------------   -------------
1/1              On               Enabled
1/2              Disabled         Disabled
1/3              Disabled         Disabled
1/4              Disabled         Disabled
...
```

Related Commands

Note

show dcb priority-flow-control interface mlag-port-channel

show dcb priority-flow-control interface mlag-port-channel <inf> [detail]

Displays DCB priority flow control configuration and status for MPO interfaces.

Syntax Description	inf	The interface number.
	detail	Adds details information to the show output.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.6000	Updated Example
Role	admin	

Example	`switch (config) # show dcb priority-flow-control interface mlag-port-channel 1 detail`
	
	`PFC: disabled`
	`Priority Enabled List:`
	`Priority Disabled List: 0 1 2 3 4 5 6 7`
	
	`PFC Port Mpo1 Information:`
	`  Port Mode  :  On`
	`  Operational state  :  Off`
	
	`No Remote Entry is Present`	

Related Commands		
Note		

show interface pfc-wd

show interface <type> <id> pfc-wd

Displays PFC watchdog information.

Syntax Description	type	Interface type: • ethernet • port-channel • mlag-port-channel
	id	Interface ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 pfc-wd`
	`1/1:`
	`--------------------------------`
	`PFC-WD admin : disable`
	`PFC-WD state :`
	`TC 0 = TC OK`
	`TC 1 = TC OK`
	`TC 2 = TC OK`
	`TC 3 = TC OK`
	`TC 4 = TC OK`
	`TC 5 = TC OK`
	`TC 6 = TC OK`
	`TC 7 = TC in DEADLOCK, duration 1 sec`	

Related Commands	pfc-wd	
Note		

5.16 Quality of Service (QoS)

5.16.1 QoS Classification

QoS classification assigns a QoS class to the packet. The QoS class of the packet is indicated internally in the switch using the switch-priority parameter (8 possible values).

Switch-priority affects the packet buffering and transmission scheduling. There are 8 possible values for switch-priority. The classification is based on the PCP and DEI fields in the VLAN tag, the DSCP field in the IP header. In addition, the default value can be configured for the incoming port. And the switch-priority of the packet also can be reconfigured by the ACL.

The switch-priority of the packet is used for priority fields re-marking at the egress.

5.16.1.1 Trust Levels

QoS classification depends on the port configuration for QoS trust level which determines which packet header fields derive the switch-priority. The following trust states are supported:

- Trust port
  - Based on port default settings
- Trust L2 (PCP,DEI)
  - Based on packet PCP,DEI fields for VLAN tagged packets
  - Else, based on the port default setting for VLAN un-tagged packets
- Trust L3 (DSCP)
  - Based on packet DSCP field for IP packets
  - Else, based on port default setting for non-IP
- Trust both
  - Based on packet DSCP for IP packets
  - Else, based on packet PCP,DEI for VLAN tagged packets
  - Else, based on the port default setting

Table 55 and figure summarize the classification rules.

Table 55 - Packet Classification Rules

Packet Type		QoS Classification Config (per Interface)			
IP/MPLS	**VLAN**	**Trust Both**	**Trust L3**	**Trust L2**	**Trust Port**
IP/MPLS	**Tagged**	DSCP	DSCP	PCP,DEI	Port Default
IP/MPLS	**Untagged**	DSCP	DSCP	Port Default	Port Default
non-IP/MPLS	**Tagged**	PCP,DEI	Port Default	PCP,DEI	Port Default
non-IP/MPLS	**Untagged**	Port Default	Port Default	Port Default	Port Default

Default switch-priority is configured as trust L2.

5.16.1.2 Switch Priority to IEEE Priority Mapping

IEEE defines priority value for a packet which is used in the switch for the pause flow control.

The device maps the switch-priority into IEEE priority value using device global switch priority to IEEE priority table.

5.16.1.3 Default QoS Configuration

Table 56 - Default QoS Configuration

Parameter	Range	Configuration
Trust level	All ports	Trust L2
DSCP to switch-priority	0-7	0
DSCP to switch-priority	8-15	1
DSCP to switch-priority	16-23	2
DSCP to switch-priority	24-31	3
DSCP to switch-priority	32-39	4
DSCP to switch-priority	40-47	5
DSCP to switch-priority	48-55	6
DSCP to switch-priority	56-63	7
PCP to switch-priority	0	0
PCP to switch-priority	1	1
PCP to switch-priority	2	2
PCP to switch-priority	3	3
PCP to switch-priority	4	4
PCP to switch-priority	5	5
PCP to switch-priority	6	6
PCP to switch-priority	7	7
Port PCP,DEI default	All ports	0
Port switch-priority when “trust port” is enabled	All ports	0
Switch-priority to IEEE priority	0	0
Switch-priority to IEEE priority	1	1
Switch-priority to IEEE priority	2	2
Switch-priority to IEEE priority	3	3
Switch-priority to IEEE priority	4	4
Switch-priority to IEEE priority	5	5
Switch-priority to IEEE priority	6	6
Switch-priority to IEEE priority	7	7

5.16.2 QoS Rewrite

Spectrum™ enables rewriting QoS identifier values (DSCP, PCP, DEI) of incoming packets.

The configuration for preserving the values or rewriting them is set per ingress port. The configuration of the new values is set per egress port and is based on the mapping from the switch-priority.

In addition, the packets that pass the router module in the switch can be configured to change the "rewrite enable" configuration as well as the switch-priority.

5.16.2.1 Switch-priority to PCP,DEI Re-marking Mapping

Packet PCP and DEI fields can be updated by the switch based on switch-priority to PCP,DEI mapping tables. The mapping can be configured per egress port.

The reason for the mapping is to enable changing interpretation between two administrative domains in the network, or when a source of data is not fully trusted, and the default values are not desired. This mapping takes effect after deriving switch-priority from the PCP,DEI fields.

5.16.2.2 Switch-priority to DSCP Re-marking Mapping

Packet DSCP field can be updated based on switch-priority to DSCP mapping tables. The mapping can be configured per egress port. MPLS packets are untouched regardless this setting.

The reason for the mapping is to enable changing interpretation between two administrative domains in the network, or when a source of data is not fully trusted. This mapping will take affect after deriving switch-priority from the DSCP field.

5.16.2.3 DSCP to Switch-priority in Router

Spectrum™ enables mapping of DSCP to switch-priority in the router using a global mapping table. This mapping has global configuration for whether to change the "Rewrite/Preserve PCP,DEI" bit. This configuration sets how the DSCP to switch-priority would affect the packet.

5.16.2.4 Default Configuration

- By default no ingress rewrite configuration is set
- By default PCP rewrite configuration in router is set
- The default mapping is as following:
  - Switch-priority=i to PCP,DEI=i,0, i=0-7
  - Switch-priority=i to DSCP=8i, i=0-7

5.16.3 Queuing and Scheduling (ETS)

After the output port of the packet is determined and the packet is buffered, it is queued for transmission. Each egress port is combined from the multi-level queuing structure. The scheduling of transmission from the queues relies on various configurations such as ETS weight, flow control, rate shaping etc.

5.16.3.1 Traffic Class

The switch-priority of the packet assigns it to a specific traffic class (TClass). The TClass of the packet determines the packet path in the queuing structure. There are 8 TCs supported by the system.

5.16.3.2 Traffic Shapers

Maximum Shapers

TCs can be configured for rate shaping as described in the following:

- TClass queues: shaper per TClass queue
- Port: shaper per port (bytes only)

Shapers support the following configurations:

- Committed Incoming Rate (CIR) [bits/packets per second]
- Committed Burst Size (CBS) [bits/packets]

Each shaper has granularity rate of 1Mb/s, 10Mb/s, 100Mb/s and 1Gb/s (or 128K, 1280K, 12M, 128M pps). The maximum CBS is 3GB or 384M packets.

Minimum Shapers

TC queues can be configured for minimal rate shaping. The minimum shaper configuration overrides all other scheduling configurations. So that if ETS or WRR scheduling allocates to a TC queue lower rate than the configured minimum, that queue receives strictly higher priority over the others. If several queues receive a rate below the configured minimum, the arbitration between them can be configured as a WRR, or as strict according to the queue index.

The configuration of min shaper is identical to the configuration of max shaper.

5.16.3.3 Default Shaper Configuration

Table 57 - Default Shaper Configuration

Parameter	Range	Configuration
Switch-priority to TC	0	0
Switch-priority to TC	1	1
Switch-priority to TC	2	2
Switch-priority to TC	3	3
Switch-priority to TC	4	4
Switch-priority to TC	5	5
Switch-priority to TC	6	6
Switch-priority to TC	7	7
Shaping	All ports	No max/min shaping configured

5.16.4 RED and ECN

Random early detection (RED) is a mechanism that randomly drops packets before the switch buffer fills up in case of congestion. Explicit congestion notification (ECN) is used for congestion control protocols (TCP and RoCE CC – DCQCN) to handle congestion before packets are dropped. RED and ECN can be configured separately or concurrently per traffic class.

Relative RED/ECN is supported on TC queues. This allows the thresholds of the drop/mark actions to behave relatively to the dynamic thresholds configured for the shared buffer.

RED/ECN drop profiles are defined according to 2 parameters as shown in Figure 26:

Figure 26: RED/ECN Drop Profiles

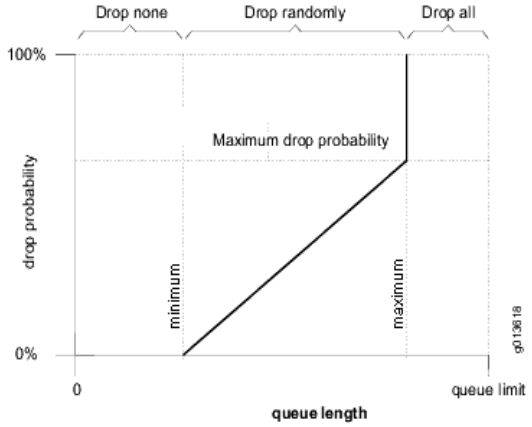

- Minimum – a threshold that defines the average queue length below which the packets are not dropped/marked
- Maximum – a threshold that defines the average queue length above which the packets are always dropped/marked

It is possible to configure the minimum and maximum thresholds to have the same value which would represent a step function from “drop none” to “drop all”.

RED/ECN is only supported for unicast traffic classes.

5.16.5 Commands

5.16.5.1 QoS Classification

vlan default priority

vlan default priority [<priority>]
no vlan default priority [<priority>]

Configures default PCP for packets arrived without VLAN tag.
The no form of the command resets the value to its default.

Syntax Description	priority Range: 0-7
Default	0
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # vlan default priority 0`
Related Commands	N/A
Notes	

vlan default dei

	vlan default dei [<dei>] **no vlan default dei [<dei>]** Configures default DEI for packets arrived without VLAN tag. The no form of the command resets the value to its default.
Syntax Description	N/A
Default	0
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # vlan default dei 0`
Related Commands	N/A
Notes	

qos trust

	qos trust [port | L2 | L3 | both] **no qos trust [port | L2 | L3 | both]** Configures QoS trust mode for the interface. The no form of the command resets the value to its default.
Syntax Description	N/A
Default	L2
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # qos trust L2`
Related Commands	N/A
Notes	

qos default switch-priority

qos default switch-priority [<switch-priority>]
no qos default switch-priority [<switch-priority>]

Configures default switch-priority for the interface when “port” trust mode is active, or for non-IP and untagged packets in other trust modes.
The no form of the command resets the value to its default.

Syntax Description	switch-priority	Range: 0-7
Default	0	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.7.00xx	Edited command definition
Role	admin	
Example	`switch (config  1/1) # qos default switch-priority 0`	
Related Commands	qos trust	
Notes		

qos map pcp dei

qos map pcp <pcp> dei <dei> [to switch-priority <switch-priority>]
no qos map pcp <pcp> dei <dei> [to switch-priority <switch-priority>]

Configures interface PCP,DEI to switch-priority mapping for IP/MPLS and non-IP/MPLS tagged packets in “L2” trust mode and for non-IP/MPLS tagged packets in “both” trust mode.
The no form of the command resets the value to its default.

Syntax Description	N/A
Default	PCP to switch-priority mapping: 0 → 0 1 → 1 2 → 2 3 → 3 4 → 4 5 → 5 6 → 6 7 → 7
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # qos map pcp 5 dei 5`
Related Commands	qos trust
Notes	

qos map dscp

qos map dscp <dscp> [to switch-priority <switch-priority>]
no qos map dscp <dscp> [to switch-priority <switch-priority>]

Configures interface DSCP to switch-priority mapping in “L3” or “both” trust mode. The no form of the command resets the value to its default.

Syntax Description	switch-priority	Range: 0-7
	dscp	Range: 0-63
Default	DSCP to switch-priority mapping:	0-7 → 0 8-15 → 1 16-23 → 2 24-31 → 3 32-39 → 4 40-47 → 5 48-55 → 6 56-63 → 7
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1) # qos map dscp 45`	
Related Commands	qos trust	
Notes		

show interfaces ethernet counters tc

show interfaces ethernet <slot/port> counters tc <priority>

Displays traffic group counters for the specified interface and priority.

Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 counters tc 3`
	
	`TC 3`
	`  0  packets`
	`  0  bytes`
	`  0  queue depth`
	`  0  unicast no buffer discard`
	`  0  WRED discard`	

Related Commands		
Note		

show interfaces ethernet counters pg

show interfaces ethernet <slot/port> counters pg <priority>

Displays port group counters for the specified interface and priority.

Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 counters pg 4`
	
	`PG 4`
	`  0  packets`
	`  0  bytes`
	`  0  queue depth`
	`  0  no buffer discard`
	`  0  shared buffer discard`	

Related Commands		
Note		

show interfaces ethernet counters pfc prio

show interfaces ethernet <slot/port> counters pfc prio <priority>

Displays priority flow control counters for the specified interface and priority.

Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Role	admin	

Example	`switch (config) # show interfaces ethernet 1/1 counters pfc prio 1`
	
	`PFC 1`
	
	`Rx`
	`  0  pause packets`
	`  0  pause duration`
	
	`Tx`
	`  0  pause packets`
	`  0  pause duration`	

Related Commands		
Note		

show qos

show qos

Displays QoS information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.1002 3.6.8008 Updated Example
Role	admin
Example	

```
switch (config) # show qos

Eth1/1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ----------------------------------------------------------------
    PCP(DEI)              DSCP                      switch-priority
    ----------------------------------------------------------------
    0(0) 0(1)             0 1 2 3 4 5 6 7           0
    1(0) 1(1)             8 9 10 11 12 13 14 15     1
    2(0) 2(1)             16 17 18 19 20 21 22 23   2
    3(0) 3(1)             24 25 26 27 28 29 30 31   3
    4(0) 4(1)             32 33 34 35 36 37 38 39   4
    5(0) 5(1)             40 41 42 43 44 45 46 47   5
    6(0) 6(1)             48 49 50 51 52 53 54 55   6
    7(0) 7(1)             56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP;
traffic-class):
    Egress Interface: Eth1/1

    ---------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ---------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
...
```

Related Commands	N/A
Notes	

show qos

show qos <port-id>

Display QoS information for Ethernet interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000 3.6.8008 Updated Example
Role	admin
Example	

```
switch (config)# show qos  1/1

Eth1/1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ----------------------------------------------------------------
    PCP(DEI)              DSCP                      switch-priority
    ----------------------------------------------------------------
    0(0) 0(1)             0 1 2 3 4 5 6 7           0
    1(0) 1(1)             8 9 10 11 12 13 14 15     1
    2(0) 2(1)             16 17 18 19 20 21 22 23   2
    3(0) 3(1)             24 25 26 27 28 29 30 31   3
    4(0) 4(1)             32 33 34 35 36 37 38 39   4
    5(0) 5(1)             40 41 42 43 44 45 46 47   5
    6(0) 6(1)             48 49 50 51 52 53 54 55   6
    7(0) 7(1)             56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):
    Egress Interface: Eth1/1

    ---------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ---------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
```

Related Commands
Note

show qos interface mlag-port-channel

show qos interface mlag-port-channel <port-id>

Display QoS information for MPO.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000 3.6.6000 Updated Example
Role	admin
Example	

```
switch (config)# show qos interface mlag-port-channel 1
Mpo1
 Trust mode: L2
 Default switch-priority: 0
 Default PCP: 0
 Default DEI: 0
 PCP,DEI rewrite: disabled
 IP PCP;DEI rewrite: enable
 DSCP rewrite: disabled

 PCP(DEI); DSCP to switch-priority mapping:
------------------------------------------------------------------
 PCP(DEI)             DSCP                      switch-priority
------------------------------------------------------------------
 0(0) 0(1)            0 1 2 3 4 5 6 7           0
 1(0) 1(1)            8 9 10 11 12 13 14 15     1
 2(0) 2(1)            16 17 18 19 20 21 22 23   2
 3(0) 3(1)            24 25 26 27 28 29 30 31   3
 4(0) 4(1)            32 33 34 35 36 37 38 39   4
 5(0) 5(1)            40 41 42 43 44 45 46 47   5
 6(0) 6(1)            48 49 50 51 52 53 54 55   6
 7(0) 7(1)            56 57 58 59 60 61 62 63   7

 PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traf-
fic-class):

 Egress Interface: Mpo1
----------------------------------------
 switch-priority   PCP(DEI)   DSCP     TC
----------------------------------------
 0                 0(0)       0        0
 1                 1(0)       8        1
 2                 2(0)       16       2
 3                 3(0)       24       3
 4                 4(0)       32       4
 5                 5(0)       40       5
 6                 6(0)       48       6
 7                 7(0)       56       7
```

Related Commands
Note

show qos interface port-channel

show qos interface port-channel <port-id>

Display QoS information for port-channel interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
	3.6.8008 Updated Example
Role	admin

Example

```
switch (config)# show qos interface port-channel 1

Po1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ------------------------------------------------------------------
    PCP(DEI)             DSCP                      switch-priority
    ------------------------------------------------------------------
    0(0) 0(1)            0 1 2 3 4 5 6 7           0
    1(0) 1(1)            8 9 10 11 12 13 14 15     1
    2(0) 2(1)            16 17 18 19 20 21 22 23   2
    3(0) 3(1)            24 25 26 27 28 29 30 31   3
    4(0) 4(1)            32 33 34 35 36 37 38 39   4
    5(0) 5(1)            40 41 42 43 44 45 46 47   5
    6(0) 6(1)            48 49 50 51 52 53 54 55   6
    7(0) 7(1)            56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):
    Egress Interface: Po1

    ----------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ----------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
```

Related Commands

Note

show qos interface l2-mapping

show qos interface <type> <port-id> l2-mapping

Displays the PCP, DEI to switch priority table.

Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	

Example	`switch (config)# show qos  1/9 l2-mapping`
	
	`PCP,DEI to switch-priority mapping:`
	`----------------------------------------`
	`PCP(DEI)  switch-priority`
	`----------------------------------------`
	`0(0) 0(1)  0`
	`1(0) 1(1)  1`
	`2(0) 2(1)  2`
	`3(0) 3(1)  3`
	`4(0) 4(1)  4`
	`5(0) 5(1)  5`
	`6(0) 6(1)  6`
	`7(0) 7(1)  7`	

Related Commands		
Note		

show qos interface l3-mapping

show qos interface <type> <port-id> l3-mapping

Displays the DSCP to switch priority table.

Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	

Example	`switch (config)# show qos  1/9 l3-mapping`
	
	`IP PCP,DEI rewrite: enabled`
	`DSCP to switch-priority mapping:`
	`-----------------------------------------`
	`DSCP                    switch-priority`
	`-----------------------------------------`
	`0 1 2 3 4 5 6 7           0`
	`8 9 10 11 12 13 14 15     1`
	`16 17 18 19 20 21 22 23   2`
	`24 25 26 27 28 29 30 31   3`
	`32 33 34 35 36 37 38 39   4`
	`40 41 42 43 44 45 46 47   5`
	`48 49 50 51 52 53 54 55   6`
	`56 57 58 59 60 61 62 63   7`	

Related Commands		
Note		

show qos interface rewrite-mapping

show qos interface <type> <port-id> rewrite-mapping

Displays the rewrite mapping of switch priority to PCP, DEI and DSCP table.

Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.8008	Updated Example
Role	admin	
Example		

```
switch (config)# show qos  1/1 rewrite-mapping

PCP,DEI rewrite   : disabled
IP PCP,DEI rewrite: enable
DSCP rewrite      : disabled

Rewrite mapping (switch-priority to PCP,DEI,DSCP):
  Egress Interface: Eth1/1

  ----------------------------------------
  switch-priority   PCP(DEI)   DSCP     TC
  ----------------------------------------
  0                 0(0)       0        0
  1                 1(0)       8        1
  2                 2(0)       16       2
  3                 3(0)       24       3
  4                 4(0)       32       4
  5                 5(0)       40       5
  6                 6(0)       48       6
  7                 7(0)       56       7
```

Related Commands	
Note	

show qos interface tc-mapping

show qos interface <type> <port-id> tc-mapping

Displays mapping from switch priority to traffic class.

Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	

Example	`switch (config)# show qos  1/9 tc-mapping`
	`Switch Priority to TC mapping:`
	`----------------------`
	`Switch Priority   TC`
	`----------------------`
	`0                 0`
	`1                 1`
	`2                 2`
	`3                 3`
	`4                 4`
	`5                 5`
	`6                 6`
	`7                 7`	

Related Commands		
Note		

show qos mapping ingress interface egress interface

show qos mapping ingress interface <type> <port-id> egress interface <type> <port-id>

Displays end to end mapping configuration: ingress to egress.

Syntax Description	type Ethernet, port-channel, or mlag-port-channel
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin
Example	

```
switch (config)# show qos mapping ingress  1/8 egress  1/9

Ingress Interface: Eth1/8
Trust mode: L2
Default Switch Priority: 0
Rewrite PCP,DEI: disabled
Rewrite DSCP   : disabled

Global Rewrite mode: enabled

PCP,DEI and DSCP to switch-priority mapping:
----------------------------------------------------------------
PCP,DEI             DSCP                      switch-priority
----------------------------------------------------------------
0(0) 0(1)           0 1 2 3 4 5 6 7           0
1(0) 1(1)           8 9 10 11 12 13 14 15     1
2(0) 2(1)           16 17 18 19 20 21 22 23   2
3(0) 3(1)           24 25 26 27 28 29 30 31   3
4(0) 4(1)           32 33 34 35 36 37 38 39   4
5(0) 5(1)           40 41 42 43 44 45 46 47   5
6(0) 6(1)           48 49 50 51 52 53 54 55   6
7(0) 7(1)           56 57 58 59 60 61 62 63   7
-----------------------------------------
 switch-priority   PCP(DEI)   DSCP     TC
-----------------------------------------
 0                 0(0)       0        0
 1                 1(0)       8        1
 2                 2(0)       16       2
 3                 3(0)       24       3
 4                 4(0)       32       4
 5                 5(0)       40       5
 6                 6(0)       48       6
 7                 7(0)       56       7
```

Related Commands	
Note	

5.16.5.2 QoS Rewrite

qos rewrite pcp

qos rewrite pcp-enable
qos rewrite pcp-disable

Enables or disables PCP,DEI rewrite on the interface.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # qos rewrite pcp-enable`
Related Commands	
Notes	

qos rewrite dscp

	qos rewrite dscp-enable **qos rewrite dscp-disable** Enables or disables DSCP rewrite on the interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # qos rewrite dscp-enable`
Related Commands	
Notes	

qos rewrite map switch-priority pcp dei

qos rewrite map switch-priority <switch-priority> pcp <pcp> dei <dei>
no qos rewrite map switch-priority <switch-priority> pcp <pcp> dei <dei>

Configures switch-priority to PCP,DEI mapping on the interface.
The no form of the command resets the value to their defaults.

Syntax Description	switch-priority	Range: 0-7
	pcp	Range: 0-7
	dei	Value: 0
Default	Switch priority to PCP,DEI mapping:	Switch priority → PCP,DEI: 0 → 0,0 1 → 1,0 2 → 2,0 3 → 3,0 4 → 4,0 5 → 5,0 6 → 6,0 7 → 7,0
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1) # qos rewrite map switch -priority 11 pcp 7 dei 0`	
Related Commands	N/A	
Notes		

qos rewrite map switch-priority dscp

	qos rewrite map switch-priority <switch-priority> dscp <dscp> **no qos rewrite map switch-priority <switch-priority> dscp <dscp>** Configures switch-priority to DSCP mapping on the interface. The no form of the command resets the value to their defaults.
Syntax Description	N/A
Default	Switch priority to DSCP mapping: 0 → 0 1 → 8 2 → 16 3 → 24 4 → 32 5 → 40 6 → 48 7 → 54
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Role	admin
Example	`switch (config  1/1) # qos rewrite map switch -priority 5 dscp 40`
Related Commands	N/A
Notes	

qos ip rewrite pcp

qos ip rewrite pcp [disable | enable | preserve]
no qos ip rewrite pcp [disable | enable | preserve]

Enables or preserves the rewrite of PCP, DEI of routed packets in egress interface. The no form of the command resets the value to their defaults.

Syntax Description	disable	No rewrite occurs
	enable	PCP,DEI are rewritten based on the mapping configured on the egress port
	preserve	Ingress interface configuration determines action
Default	Default is “preserve” when router is disabled Default is “enable” when router is enabled (Router can be enabled/disabled using the “ip routing” command)	
Configuration Mode	config	
History	3.6.1002	
Role	admin	
Example	`switch (config) # qos ip rewrite pcp enable`	
Related Commands	N/A	
Notes		

show qos ip rewrite

	show qos ip rewrite Displays configuration of the rewrite of PCP, DEI of routed packets in egress interface
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000
Role	admin
Example	`switch (config)# show qos ip rewrite` `IP rewrite PCP: enable`
Related Commands	qos ip rewrite pcp
Notes	

5.16.5.3 Queuing and Scheduling (ETS)

dcb ets enable

	dcb ets enable **no dcb ets enable** Sets the switch egress scheduling mode to be weighted round robin. The no form of the command sets the switch egress scheduling mode to be strict priority.
Syntax Description	N/A
Default	ETS is enabled
Configuration Mode	config
History	3.1.0000 3.6.1002 Updated Note
Role	admin
Example	`switch (config)# dcb ets enable`
Related Commands	show dcb ets
Note	

bind switch-priority

bind switch-priority [<priority_1> [<priority_2] .. <priority_n>]]
no bind switch-priority [<priority>]

Configures binding of switch-priority to traffic class.
The no form of the command:

- When run in the interface configuration mode: Resets to default the binding of all switch-priorities from all traffic classes
- When run in the interface's traffic class: Negates the binding of a specific switch-priority from a specific traffic class

Syntax Description	N/A
Default	Switch priority to traffic class mapping: 0 → 0 1 → 1 2 → 2 3 → 3 4 → 4 5 → 5 6 → 6 7 → 7
Configuration Mode	config interface ethernet config interface ethernet traffic-class config interface port-channel config interface port-channel traffic-class config interface mlag-port-channel config interface mlag-port-channel traffic class
History	3.6.1002
Role	admin
Example	`switch (config  1/1 traffic-class 0) # bind switch-property 1`
Related Commands	N/A
Notes	• Context is egress interface traffic class

bandwidth guaranteed

bandwidth guaranteed [<rate>]
no bandwidth guaranteed [<rate>]

Configures the minimum bandwidth for outbound traffic.

Syntax Description	rate	Rate in GbE Range: 0 - max speed supported
Default	0	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 traffic-class 0) # bandwidth guaranteed 0.4G`	
Related Commands	N/A	
Notes	• Context is egress interface traffic class • Bandwidth guaranteed rate determines the bandwidth guaranteed by the switch for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G	

bandwidth shape

bandwidth shape [<rate>]
no bandwidth shape [<rate>]

Configures the bandwidth shaper for outbound traffic.

Syntax Description	rate	Rate in GbE Range: 0 - max speed supported
Default	Maximum port rate (100GbE on Spectrum™ based switches)	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 traffic-class 7) # bandwidth shape 0.4G`	
Related Commands	N/A	
Notes	• Context is egress interface traffic class and/or port • Bandwidth shape rate determines the bandwidth of the shaper for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G	

dcb ets

	dcb ets [strict | wrr <weight>] **no dcb ets [strict | wrr <weight>]** Configures ETS mode to strict or WRR.
Syntax Description	weight
Default	Default is WRR with the following default weights. Traffic class to weight mapping: 0 → 12 1 → 13 2 → 12 3 → 13 4 → 12 5 → 13 6 → 12 7 → 13
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class
History	3.6.1002
Role	admin
Example	`switch (config  1/1 traffic-class 1) # dcb ets wrr 50`
Related Commands	N/A
Notes	Context is egress interface traffic class

show dcb ets

show dcb ets [interface {ethernet | mlag-port-channel | port-channel} <if-id>]

Displays ETS information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.1002 3.6.5000 Updated Example
Role	admin

Example	`switch (config)# show dcb ets  1/1`
	`Eth1/1:`
	`Interface Bandwidth Shape [Mbps]: N/A`
	`Multicast unaware mapping: disabled`
	
	`Flags:`
	`S.Mode: Scheduling Mode [Strict/WRR]`
	`D: -`
	`W: Weight`
	`Bw.Sh: Bandwidth Shaper`
	`Bw.Gr: Bandwidth Guaranteed`
	
	`ETS per TC:`
	`----------------------------------------------------------------`
	`TC  S.Mode  W  W(%)  BW Sh.(Mbps)  BW Gr.(Mbps)`
	`----------------------------------------------------------------`
	`0  WRR  12  12  N/A  0`
	`1  WRR  13  13  N/A  0`
	`2  WRR  12  12  N/A  0`
	`3  WRR  13  13  N/A  0`
	`4  WRR  12  12  N/A  0`
	`5  WRR  13  13  N/A  0`
	`6  WRR  12  12  N/A  0`
	`7  WRR  13  13  N/A  0`

Related Commands	N/A
Notes	

5.16.5.4 RED & ECN

traffic-class congestion-control

traffic-class <tc> congestion-control [red | ecn | both] [minimum- absolute <min> maximum-absolute <max> | minimum-relative <min> maximum-relative <max>]
no traffic-class <tc> congestion-control

Enables RED/ECN marking for traffic class queue.
The no form of the command disables RED/ECN marking for traffic class queue.

Syntax Description	tc	Traffic class. Range: 0-7.
	red	Enables random early detection for traffic class queue
	ecn	Enables explicit congestion notification for traffic class queue
	both	Enables both RED and ECN marking for traffic class queue
	minimum-absolute	Set minimum-absolute value (in KBs) for marking traffic-class queue
	maximum-absolute	Set maximum-absolute value (in KBs) for marking traffic-class queue
	minimum-relative	Set minimum-relative value (in percentage) for marking traffic-class queue
	maximum-relative	Set maximum-relative value (in percentage) for marking traffic-class queue
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.5.1000	
Role	admin	
Example	`switch (config interfaces ethernet 1/1)# traffic-class 0 congestion-control both minimum-relative 50 maximum-relative 80`	
Related Commands		
Note		

show interfaces ethernet congestion-control

show interfaces ethernet congestion-control

Displays specific interface congestion control information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.1000
Role	admin
Example	

```
switch (config)# show interfaces ethernet 1/1 congestion-control
: 1/1

ECN marked packets: 0
TC-0
        Mode: ECN
        Threshold mode: absolute
        Minimum threshold: 0 KB
        Maximum threshold: 200 KB
        RED dropped packets: 0
TC-1
        Mode: RED
        Threshold mode: relative
        Minimum threshold: 0%
        Maximum threshold: 100%
        RED dropped packets: 0
TC-2
        Mode: none
TC-3
        Mode: none
TC-4
        Mode: ECN
        Threshold mode: relative
        Minimum threshold: 25%
        Maximum threshold: 80%
        RED dropped packets: 0
TC-5
        Mode: none
TC-6
        Mode: both
        Threshold mode: absolute
        Minimum threshold: 100 KB
        Maximum threshold: 200 KB
        RED dropped packets: 0
TC-7
        Mode: none

switch (config) #
```

Related Commands
Note

bind switch-priority

bind switch-priority [<priority_1> [<priority_2] .. <priority_n>]]
no bind switch-priority [<priority>]

Configures binding of switch-priority to traffic class.
The no form of the command:

- When run in the interface configuration mode: Resets to default the binding of all switch-priorities from all traffic classes
- When run in the interface's traffic class: Negates the binding of a specific switch-priority from a specific traffic class

Syntax Description	N/A
Default	Switch priority to traffic class mapping: 0 → 0 1 → 1 2 → 2 3 → 3 4 → 4 5 → 5 6 → 6 7 → 7
Configuration Mode	config interface ethernet config interface ethernet traffic-class config interface port-channel config interface port-channel traffic-class config interface mlag-port-channel config interface mlag-port-channel traffic class
History	3.6.1002
Role	admin
Example	`switch (config  1/1 traffic-class 0) # bind switch-property 1`
Related Commands	N/A
Notes	Context is egress interface traffic class

bandwidth guaranteed

bandwidth guaranteed [<rate>]
no bandwidth guaranteed [<rate>]

Configures the minimum bandwidth for outbound traffic.

Syntax Description	rate	Rate in GbE Range: 0 - max speed supported
Default	0	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 traffic-class 0) # bandwidth guaranteed 0.4G`	
Related Commands	N/A	
Notes	• Context is egress interface traffic class • Bandwidth guaranteed rate determines the bandwidth guaranteed by the switch for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G	

bandwidth shape

bandwidth shape [<rate>]
no bandwidth shape [<rate>]

Configures the bandwidth shaper for outbound traffic.

Syntax Description	rate	Rate in GbE Range: 0 - max speed supported
Default	Maximum port rate (100GbE on Spectrum™ based switches)	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 traffic-class 7) # bandwidth shape 0.4G`	
Related Commands	N/A	
Notes	• Context is egress interface traffic class and/or port • Bandwidth shape rate determines the bandwidth of the shaper for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G	

dcb ets

	dcb ets [strict | wrr <weight>] **no dcb ets [strict | wrr <weight>]** Configures ETS mode to strict or WRR.
Syntax Description	weight
Default	Default is WRR with the following default weights. Traffic class to weight mapping: 0 → 12 1 → 13 2 → 12 3 → 13 4 → 12 5 → 13 6 → 12 7 → 13
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class
History	3.6.1002
Role	admin
Example	`switch (config  1/1 traffic-class 1) # dcb ets wrr 50`
Related Commands	N/A
Notes	Context is egress interface traffic class

show dcb ets

show dcb ets [interface {ethernet | mlag-port-channel | port-channel} <number>]

Displays ETS information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.1002 3.6.5000 Updated example output
Role	admin
Example	

```
switch (config)# show dcb ets  1/1
Eth1/1:
Interface Bandwidth Shape [Mbps]: N/A
Multicast unaware mapping: disabled

Flags:
   S.Mode: Scheduling Mode [Strict/WRR]
   D: -
   W: Weight
   Bw.Sh: Bandwidth Shaper
   Bw.Gr: Bandwidth Guaranteed

ETS per TC:
------------------------------------------------------------------
TC   S.Mode      W      W(%)     BW Sh.(Mbps)    BW Gr.(Mbps)
------------------------------------------------------------------
0    WRR         12     12       N/A             0
1    WRR         13     13       N/A             0
2    WRR         12     12       N/A             0
3    WRR         13     13       N/A             0
4    WRR         12     12       N/A             0
5    WRR         13     13       N/A             0
6    WRR         12     12       N/A             0
7    WRR         13     13       N/A             0
```

Related Commands	N/A
Notes	

traffic-class congestion-control

traffic-class <tc> congestion-control [red | ecn | both] [minimum- absolute <min> maximum-absolute <max> | minimum-relative <min> maximum-relative <max>]
no traffic-class <tc> congestion-control

Enables RED/ECN marking for traffic class queue.
The no form of the command disables RED/ECN marking for traffic class queue.

Syntax Description	tc	Traffic class. Range: 0-7.
	red	Enables random early detection for traffic class queue
	ecn	Enables explicit congestion notification for traffic class queue
	both	Enables both RED and ECN marking for traffic class queue
	minimum-absolute	Set minimum-absolute value (in KBs) for marking traffic-class queue
	maximum-absolute	Set minimum-absolute value (in KBs) for marking traffic-class queue
	minimum-relative	Set minimum-relative value (in percentage) for marking traffic-class queue
	maximum-relative	Set maximum-relative value (in percentage) for marking traffic-class queue
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.5.1000	
Role	admin	
Example	`switch (config interfaces ethernet 1/1)# traffic-class 0 congestion-control both minimum-relative 50 maximum-relative 80`	
Related Commands		
Note		

show interfaces ethernet congestion-control

show interfaces ethernet congestion-control

Displays specific interface congestion control information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.1000
Role	admin

Example	`switch (config)# show interfaces ethernet 1/1 congestion-control`
	`: 1/1`
	
	`ECN marked packets: 0`
	`TC-0`
	`Mode: ECN`
	`Threshold mode: absolute`
	`Minimum threshold: 0 KB`
	`Maximum threshold: 200 KB`
	`RED dropped packets: 0`
	`TC-1`
	`Mode: RED`
	`Threshold mode: relative`
	`Minimum threshold: 0%`
	`Maximum threshold: 100%`
	`RED dropped packets: 0`
	`TC-2`
	`Mode: none`
	`TC-3`
	`Mode: none`
	`TC-4`
	`Mode: ECN`
	`Threshold mode: relative`
	`Minimum threshold: 25%`
	`Maximum threshold: 80%`
	`RED dropped packets: 0`
	`TC-5`
	`Mode: none`
	`TC-6`
	`Mode: both`
	`Threshold mode: absolute`
	`Minimum threshold: 100 KB`
	`Maximum threshold: 200 KB`
	`RED dropped packets: 0`
	`TC-7`
	`Mode: none`
	
	`switch (config) #`

Related Commands
Note

5.17 Shared Buffers

All successfully received packets by a switch are stored on internal memory from the time they are received until the time they are transmitted. The packet buffer is fully shared between all physical ports and is hence called a shared buffer. Buffer configuration is applied in order to provide lossless services and to ensure fairness between the ports and priorities.

The buffer mechanism allows defining reserved memory allocation and limiting the usage of memory based on incoming/outgoing ports and priority of the packet. In addition, the buffer can be divided into static pools, each for a specific set of priorities. Buffer configuration mechanism allows fair enforcement from both ingress and egress sides.

The standard configuration mode allows a simple and concise configuration manner by hiding direct buffer access from user, and collecting all the required configuration settings into “traffic pools”. Users that wish to gain full control of entire buffers set can do so by enabling advanced buffer configuration.

5.17.1 Traffic Pool Configuration

The set of configurations which will obtain the optimal shared buffer behavior according to user requirements can be applied by dividing priorities into “traffic pools”. A traffic pool is a logical representation of a traffic profile instance which is supposed to handle all buffer related allocation on the ingress and egress sides to allow fluent flow of the traffic.

Available traffic pool types are as follows:

- Lossy – for standard lossy traffic. This is the default type for all traffic.
- Lossless – for traffic which cannot suffer any loss. Using this type enables a flow control mechanism for the mapped priority as well as setting headroom and Xon/Xoff parameters for the relevant ingress PG buffer.
- Lossy-MC – for layer 2 multicast traffic which requires special care due to stream duplication on the egress side over several ports.

There is no restriction for priority mapping to traffic pools. User can map all priorities to a single traffic pool or create a separate traffic pool for each priority. By default, all memory will be equally divided between all active traffic pools. User can set a memory percentage for a traffic pool out of the entire shared buffer. A state of over-subscription (where sum of percentage is bigger than 100%) is admissible although not advised.

A traffic pool will become functional if at least one priority is mapped to it. Each functional traffic pool will be matched by an iPool, ePool and iPort.PG buffer on each interface. For further detail see Section 5.17.3, “Advanced Buffer Configuration,” on page 1107.

5.17.2 Lossless Traffic

5.17.2.1 Priority-flow-control

Enabling lossless traffic flow requires relevant switch-priority (see Packet Classification) to be mapped to a traffic pool type “Lossless”. This could be applied through one of the following methods:

- Create a new custom lossless traffic pool, and map the switch-priority to the newly created traffic pool. In this case, PFC configuration is automatic.
  Example:

```
switch (config) # traffic pool my_pool type lossless
switch (config) # traffic pool my_pool map switch-priority 0
```

- Enabling DCB PFC over the said switch-priority along with enabling DCB PFC globally. This will result in mapping of the priority to the lossless-default traffic pool which is reserved merely for this purpose. In addition it is required to enable DCB PFC for the relevant interfaces as well.

When setting lossless traffic configuration, it is strongly recommended to stick with one of the upper modes rather than a combination of them.

5.17.2.2 Flowcontrol (Global pause)

Utilizing global pause mechanism requires “flowcontrol” to be enabled over the desired port and the port's default priority must be set to switch-priority 3 to configure lossless traffic over the port. The configuration steps are described in Section 5.17.2.1, “Priority-flow-control,” on page 1106.

To ensure all incoming packets are subjected to the global pause mechanism, the port's trust mode must be set to “port”.

Example:

```
switch (config)# traffic pool my_pool type lossless
switch (config)# traffic pool my_pool map switch-priority 3
switch (config)#  1/1 flowcontrol send on force
switch (config)#  1/1 flowcontrol receive on force
switch (config)#  1/1 qos default switch-priority 3
switch (config)#  1/1 qos trust port
```

5.17.3 Advanced Buffer Configuration

5.17.3.1 Packet Buffering Classification

When a packet arrives to the switch it is classified according to its ingress port, egress port, and layer 2 and layer 3 header fields. The following terms are used to handle packet classification within the switch:

- Port
  - Ingress port (iPort) – the port which the packet is received on
  - Egress port (ePort) – the port which the packet is transmitted on
- Pool
  - Ingress pool (iPool) – the memory pool on which the packet is counted on the ingress side
  - Egress pool (ePool) – the memory pool on which the packet is counted on the egress side
- Priority
  - Switch priority (SP) – internal identifier of the packet priority which is used as a key for several internal switch functions and decisions, including buffering. The SP of the packet

is assigned according to a port's trust level configuration and packet QoS identifiers in the header (PCP, DEI, DSCP).

- Priority group (PG) – PG is combined of a group of SPs. It is used for grouping packets of several switch priorities into a single ingress buffer space. PG range is from 0-7, while PG 9 is reserved for control traffic.
- Traffic class (TC) – TC is combined of a group of SPs. It is used for grouping packets of several switch priorities into a single egress queue and buffer space. TC range is from 0-15, while TC 8-15 is reserved for multicast traffic and TC 16 is reserved for control traffic.

Buffer configuration mechanism provides a way to allocate buffer space for specific traffic types by configuring buffers of the following types.

- iPort.PG – traffic which arrives on a specific port and is mapped to a specific PG
- iPort (iPort.pool) – traffic which arrives on a specific port and is counted on a specific iPool. This sums all iPort.PG mapped to the said iPool.
- ePort.TC – traffic which is transmitted on a specific port and mapped to a specific TC
- ePort (ePort.pool) – traffic which is transmitted on a specific port and counted on a specific ePool. It should sum up all ePort.TCs mapped to the said ePool.

Since multicast packets are duplicated among egress ports, to allow consistent packet counting on ingress and egress sides, the following buffers types are used:

- MC.SP – multicast traffic which is classified per specific switch-priority. Counting occurs on egress side prior to packet duplication.
- ePort.mc – multicast traffic which is going to be transmitted on a specific port

5.17.3.2 Buffer Allocation

For the aforementioned classification parameters, a buffering region can be allocated. The buffering region is defined as a set of one of the following: {iPort}, {iPort.pg}, {ePort}, {ePort.TC}, {MC} or {MC.SP}.

For buffer regions, reserved and shared buffering quotas are allocated based on the following configuration parameters:

- Reserved allocation (size) – guaranteed buffering quota for the region which is not shared with other regions
- Shared allocation (shared) – best-effort buffering quota for the region which can be shared with other regions and allocated dynamically. Region usage cannot overflow this quota. Shared allocation can be set using static or dynamic threshold.
- Shared pool – static bound from which the shared space is dynamically allocated

The iPort.PG buffer can be configured to work in one of two modes:

- Lossy – for lossy traffic
- Lossless – for lossless traffic. In this mode, the user must define the flow control thresholds (Xoff, Xon). Reaching Xoff threshold in a PG buffer occupancy will generate "pause" frames to the sender. Reaching Xon threshold ceases "pause" frames transmission. The reserved allocation for this buffer should be at least the value of Xoff to allow sufficient ingress packet buffering for applying Xon/Xoff thresholds.

After initial admittance to headroom buffer—in which its egress port, TC, and ingress PG are defined—a packet is evaluated for eligibility for being stored in the buffer space until it is forwarded. Buffer eligibility is defined based on the following conditions:

1. If current usage is below allocation thresholds for all four shared:
   - iPort.PG && iPort && ePort.TC && ePort
2. If there is available quota within at least one of the four reserved allocation regions:
   - For lossy traffic: iPort.PG || iPort || ePort.TC || ePort
   - For lossless traffic: ePort.TC || ePort. Ingress check is not performed since all the ingress reserved space is allocated for headroom.

If a packet is not eligible for buffering:

- For lossy traffic: Packet is dropped
- For lossless traffic: Packet stays in headroom on which Xon/Xoff thresholds are applied

5.17.3.3 Pools

Shared buffer space can be statically divided among multiple pools on the ingress side (iPools) and the egress side (ePools). Each buffer is a region that is mapped to a specific pool.

Each pool has the following parameters:

- Size – the total size which is shared among the regions allocated to that pool. The pool's size binds the amount of cumulative shared usage of the regions that are mapped to the pool. The size can be set to infinite value, in which case occupancy of this pool will not be taken into consideration upon admittance of the packet.

Note: The pool size does not include the reserved sizes of regions.

- Mode – working mode
  - Static – each region has a static maximum threshold defined in bytes. The user sets the maximum shared quota for this buffer from a specific pool by providing a percentage out of the bounded pool size. If the size is set to infinite, shared quota for mapped buffers gets set in bytes.
  - Dynamic – each region has a dynamic maximal threshold defined as alpha (α) which is the ratio between the current region usage and the pool's free space (equal to the pool usage subtracted from pool size):
    - α accepts the following values 0, 1/128, 1/64, …1/2,1,2,…,64, infinity
    - Buffer acceptance condition is: region_usage < α*free pool space

The port region is counted against the pool to which the PG/TC region of the packet is mapped.

5.17.3.4 Usage Counting

A packet is counted once on the ingress side and on the egress side.

Table 58 - Single Packet Usage Counting

Direction	Traffic Type	Counting Buffers
Ingress		iPort.PG, iPort
Egress	Unicast	ePort.TC, ePort
	Multicast	MC.SP, ePort.mc

5.17.3.5 Control Traffic Buffering

Control packets are buffered in dedicated pools: iPoolCtrl, ePoolCtrl. Furthermore, each port has a set of buffers which are dedicated to control:

- iPort: iPort.iPoolCtrl
- iPort.PG: iPort.pg9
- ePort: ePort.ePoolCtrl
- ePort.TC: iPort.tc16

All control buffers are mapped to control pools and are not configurable.

5.17.3.6 Default Configuration

The default, out-of-box configuration provides the following settings:

Pools:

- iPool0, ePool0 – default pools for all data traffic. Set to dynamic mode with size of the entire shared buffer each.
- iPoolCtrl, ePoolCtrl – dynamic pools dedicated for control with size of 256KB each
- ePool15 – multicast pool with static mode and infinite size

Buffers:

- All buffer configuration (apart from MC.SP) is similar for all ports
- All switch-priorities are mapped to PG0
- Each switch-priority is mapped to a corresponding TC buffer (i-to-i)

Buffer	Reserved	Shared [%/α/Byte]	Pool	Comment
iPort.iPool0	10KB	alpha 8	iPool0 (fixed)	
iPort.iPoolCtrl	0	alpha 8	iPoolCtrl	iPort control buffer
iPort.pg0	0 (20KB headroom)	alpha 8	iPool0	
iPort.pg9	10KB	alpha 8	iPoolCtrl	iPort.pg control buffer
ePort.ePool0	10KB	alpha 8	ePool0 (fixed)	
ePort.ePoolCtrl	0	alpha 8	ePoolCtrl	ePort control buffer

Buffer	Reserved	Shared [%/α/Byte]	Pool	Comment
ePort.mc	10KB	90KB	ePool15 (fixed)	Multicast
ePort.tc0-7	1KB	alpha 8	ePool0	
ePort.tc16	1KB	alpha 8	ePoolCtrl	ePort.tc control buffer
MC.SP0-7	0	alpha ¼	ePool0	Global multicast

5.17.3.7 Configuration Example

The following example exhibits how to divide the buffer among traffic priorities in advanced buffer management mode. Assuming that over an out-of-box lossy default configuration is set, the user here configures buffering for lossless traffic classified to switch-priority 1, over Ethernet interfaces 1/1 and 1/5.

The changes on the default configuration are as follows:

- Advanced buffer management is enabled
- Ingress:
  - iPool1 is assigned a size of 13MB
  - Switch-priority is bound to PG1 to allow separate configuration settings
  - PG1 is mapped to selected pool iPool1, classified as lossless and set sufficient headroom (reserved size) of 85KB. Xon/Xoff thresholds are set to 20KB. The shared alpha coefficient is set to 1.
  - iPort.pool1 buffer receives reserved size of 10k and shared coefficient of alpha 1.
- Egress:
  - ePool1 is assigned an infinite size according to recommended lossless traffic settings
  - TC1 (to which switch-priority is mapped by default) is mapped to the selected pool ePool1, and receives reserved size 0 and an infinite shared threshold
  - ePort.mc buffer receives reserved size 0 and an infinite shared threshold
  - ePort.pool1 buffer receives reserved size 0 and an infinite shared threshold
  - MC.SP1 buffer is mapped to egress pool ePool1, and gets reserved size 0 and an infinite shared threshold
- Finally, priority-flow-control is enabled over switch-priority 1, and over the selected ports.

Example:

```
switch (config) # advanced buffer management force
# Pool configuration
switch (config) # pool iPool1 size 13680063 type dynamic
switch (config) # pool ePool1 size inf type static
# Ingress buffer configuration
switch (config) #  1/1 ingress-buffer iPort
 pool iPool1 reserved 10k shared alpha 1
```

```
switch (config) #  1/1 ingress-buffer iPort.pg1
 bind switch-priority 1
switch (config) #  1/1 ingress-buffer iPort.pg1
 map pool iPool1 type lossless reserved 85k xoff 20k xon 20k shared alpha 1
switch (config) #  1/1 egress-buffer ePort
 pool ePool1 reserved 0 shared size inf
switch (config) #  1/1 egress-buffer ePort.tc1
 map pool ePool1 reserved 0 shared size inf
switch (config) #  1/1 egress-buffer ePort.mc
 reserved 0 shared size inf
# Egress buffer configuration
switch (config) #  1/5 ingress-buffer iPort
 pool iPool1 reserved 10k shared alpha 1
switch (config) #  1/5 ingress-buffer iPort.pg1
 bind switch-priority 1
switch (config) #  1/5 ingress-buffer iPort.pg1
 map pool iPool1 type lossless reserved 85k xoff 20k xon 20k shared alpha 1
switch (config) #  1/5 egress-buffer ePort pool
 ePool1 reserved 0 shared size inf
switch (config) #  1/5 egress-buffer ePort.tc1
 map pool ePool1 reserved 0 shared size inf
switch (config) #  1/5 egress-buffer ePort.mc
 reserved 0 shared size inf
# MC buffer configuration
switch (config) # pool ePool1 mc-buffer mc.sp1 reserved 0 shared size inf
# PFC configuration
switch (config) # dcb priority-flow-control enable force
switch (config) # dcb priority-flow-control priority 1 enable
switch (config) #  1/1 dcb priority-flow-control mode on
switch (config) #  1/5 dcb priority-flow-control mode on
```

5.17.3.8 Exceptions to Legal Shared Buffer Configuration

The following configurations are permissible in spite of them not being logical since they are useful to the user in specific advanced situations:

- Global scenarios:
  - Traffic pool memory over-subscription (total X%) and Traffic pools with size ‘Auto’ are not allocated.

    In this scenario, two or more traffic pools are configured so the sum of their sizes (specified in the percentage units) is more than 100%. In this case, upon high utilization, traffic “fights” for resources (free pool memory) and can be lost.
  - Switch priority X is mapped to a non-lossless traffic pool, but PFC is enabled on it, or

    Switch priorities X-1,X are mapped to a non-lossless traffic pool, but PFC is enabled on them.

    In these scenarios, switch priority X is mapped to a lossy or lossy-MC traffic pool (traffic is not important and traffic loss is allowed), but pause packet generation (PFC) also is enabled over this priority. These cases are allowed if the user expects traffic to be dropped but has enabled PFC to prevent it.

- Switch priority X is mapped to a lossless traffic pool, but PFC is disabled on it, or

  Switch priorities X-1,X are mapped to a lossless traffic pool, but PFC is disabled on them.

  As opposed to the previous scenarios, here the traffic pool is created as lossless, but pause packet generation is disabled. In these cases, the user expects traffic not to have drops, but it can be dropped.

- Per interface scenarios:
  - <if-id> TC X is mapped to more than one traffic pool, or TCs X,X+1 are mapped to more than one traffic pool.

    In these scenarios, traffic class buffers share the same switch priority and are mapped to two different traffic pool. In this cases, with different traffic pool configuration, behavior of traffic is not determined.
  - <if-id> switch priority X is lossless but neither PFC nor FC is not enabled on this interface, or

    Switch priorities X-1,X are lossless but neither PFC nor FC is enabled on this interface.

    In these scenarios, the user has created a lossless traffic pool and expects that traffic would not be dropped, but pause packet generation (PFC and FC) is disabled on the interface. In these cases, traffic can be dropped.
  - <if-id> has FC enabled, but default priority 0 is not mapped to lossless traffic pool and FC may not be functional.

    In this scenario, global pause packet (FC) generation is enabled on the interface, but default switch priority (traffic arriving to the switch without priority tagging is assigned the default switch priority) is not in lossless traffic pool. In this case, traffic cam be dropped.
  - <if-id> has insufficient headroom allocation to fulfill configuration derived requirements (MTU, speed, cable-length).

    In this scenario, combination of MTU, speed, cable-length, and amount of lossless traffic pools consumes all free headroom memory. In this case, not all required buffers are configured correctly and traffic can be dropped.

5.17.4 Commands

traffic pool

traffic pool <name> [force]
no traffic pool <name> [force]

Creates a traffic pool and enters the traffic pool context on prefix mode enabled. The no form of the command deletes a traffic pool.

Syntax Description	name	String up to 20 characters
	force	Enforces configuration
Default	N/A	
Configuration Mode	config	
History	3.6.5000	
Role	admin	
Example	`switch (config)# traffic pool name` `switch (config pool name)#`	
Related Commands		
Note		

type

	type <type> **no type <type>** Configures the traffic pool type. The no form of the command resets a traffic pool.	
Syntax Description	type	• lossless • lossy • lossy-mc
Default	Lossy	
Configuration Mode	config pool	
History	3.6.5000	
Role	admin	
Example	`switch (config pool name)# type lossless`	
Related Commands		
Note	When using “traffic pool <name> type <type>”, if the traffic pool does not exist then it is created.	

map switch-priority

map switch-priority <list-of-priorities>
no map switch-priority <list-of-priorities>

Maps switch-priorities to the traffic pool.
The no form of the command unmaps switch-priorities.

Syntax Description	list-of-priorities Range: 0-7
Default	N/A
Configuration Mode	config pool
History	3.6.5000
Role	admin
Example	`switch (config pool name)# map switch-priority 2 3 1 7`
Related Commands	
Note	When using “traffic pool <name> map switch-priority <list-of-priorities>”, if the traffic pool does not exist then it is created.

type map switch-priority

type {lossless | lossy | lossy-mc} map switch-priority <priority>
no type {lossless | lossy | lossy-mc} map switch-priority

Configures type of traffic pool and maps switch-priorities to it.
The no form of the command unmaps switch-priorities.

Syntax Description	type	• lossless • lossy • lossy-mc
	priority	Range: 0-7
Default	Type: Lossy	
Configuration Mode	config pool	
History	3.6.5000	
Role	admin	
Example	`switch (config pool name)# type lossy-mc map switch-priority 2 3 1 7`	
Related Commands		
Note	When using “traffic pool <name> type <type> map switch-priority <priority>”, if the traffic pool does not exist the it is created.	

memory percent

memory percent [<percent>]
no memory percent [<percent>]

Sets traffic pool size in percentage out of entire shared buffer memory.
The no form of the command resets this parameter to its default.

Syntax Description	percent Range: 0.00-100.00 or “auto”
Default	Auto
Configuration Mode	config pool
History	3.6.5000
Role	admin
Example	`switch (config pool name)# memory percent 50.03`
Related Commands	
Note	• Setting “auto” value ensures fair memory division between all traffic pools with “auto” size • Over-subscription of more than 100% is allowed but not recommended, and causes an exception to be displayed in the “Exceptions list” in “show traffic pool” command output. See Section 5.17.3.8, “Exceptions to Legal Shared Buffer Configuration,” on page 1112 for more details.

advanced buffer management

advanced buffer management [force]
no advanced buffer management [force]

Enable the advanced mode shared buffer configuration.
The no form of the command disables the advanced mode shared buffer configuration.

Syntax Description	force	Run command skipping confirmation prompt
Default	Disabled	
Configuration Mode	config	
History	3.6.5000	
	3.6.8008	Updated Note field
Role	admin	
Example	`switch (config)# advanced buffer management force` `This will reset all configuration to default. Type 'yes' to confirm: yes`	
Related Commands		
Note	When moving advanced buffer management from disable to enable, buffer/PFC configuration returns all shared buffer configuration to default.	

ingress-buffer

ingress-buffer <buffer-name>
no ingress-buffer <buffer-name>

Creates and enters the ingress buffer context.
The no form of the command deletes an existing buffer.

Syntax Description	buffer-name	Name of ingress buffer
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1)# ingress-buffer iPort.pg1` `switch (config  1/1 ingress-buffer iPort.pg1)#`	
Related Commands		
Note	iPort.pg9 is reserved for control traffic and hence cannot be edited	

egress-buffer

egress-buffer <buffer-name>
no egress-buffer <buffer-name>

Creates and enters the buffer context.
The no form of the command deletes an existing buffer.

Syntax Description	buffer-name	Name of egress buffer
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1)# egress-buffer ePort.tc4` `switch (config  1/1 egress-buffer ePort.tc4)#`	
Related Commands		
Note	ePort.tc16 is reserved for control traffic and hence cannot be edited	

reserved shared size

reserved <value> shared size <size>
no reserved <value>

Configures the ePort.mc multicast-buffer.
The no form of the command resets buffer to default configuration.

Syntax Description	buffer-name	Name of egress buffer
	value	Amount of reserved memory for buffer in bytes
	shared size	Shared memory in bytes or “infinite”
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/1 egress-buffer ePort.mc)# reserved 5k shared alpha 1/128`	
Related Commands		
Note	• ePort.tc16 is reserved for control traffic and hence cannot be edited • It is possible to use “K” and “M” to define shared size	

pool size type

pool <pool-name> size <value> type {static | dynamic}
no pool <pool-name> size <value> type {static | dynamic}

Creates pool.
The no form of the command deletes pool.

Syntax Description	pool-name	Possible values: • ePool0 ... ePool6 • iPool0 ... iPool6
	size	Size of pool in bytes, or “inf” for infinite
Default	According to system default OOB configuration	
Configuration Mode	config	
History	3.6.5000	
Role	admin	
Example	`switch (config)# pool iPool2 size 2M type dynamic` `switch (config)# pool iPool2 size static type static`	
Related Commands		
Note	It is possible to use “K” for kilobytes and “M” for megabytes to define pool size.	

pool reserved shared

pool <pool-name> reserved <reserved> shared <shared units> <shared>
no pool <pool-name>

Configures the buffer.
The no form of the command resets the values to their default.

Syntax Description	pool-name	Possible values: iPool0-iPool7
	reserved	Amount of reserved memory for the buffer in bytes
	shared units	The amount of shared memory for this buffer Possible values: alpha, max, size • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 … 1, 2, 4, … 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 ingress-buffer iPort)# pool iPool0 reserved 90K shared alpha 1/8`	
Related Commands		
Note		

map pool type reserved

map [pool <pool name> type <type> [xoff <xoff-value> xon <xon value>] reserved <reserved size> shared <shared units> <shared size>]

Maps iPort.pg buffer to a given pool and sets its reserved and shared sizes.
The no form of the command resets buffer to default pool mapping and configuration.

Syntax Description	pool-name	Possible values: iPool0 … iPool7
	type	Possible values: lossy, lossless
	reserved size	Amount of reserved memory for the buffer in bytes
	shared units	Possible values: size, alpha, max
	shared size	The amount of shared memory for this buffer • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 … 1, 2, 4, … 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite Shared size depends on type and size of the given pool: • For static pool shared size is in packets • For dynamic pool shared size is in alpha units • For static pool with infinite size only alpha infinite is supported
	xoff	Relevant only on lossless type, Xoff threshold in bytes
	xon	Relevant only on lossless type, Xon threshold in bytes
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet ingress-buffer	
History	3.6.1002	
	3.6.5000	Updated command syntax
Role	admin	
Example	`switch (config  1/9 ingress-buffer iPort.pg5)# map pool iPool6 type lossy reserved 3k shared alpha 2` `switch (config  1/9 ingress-buffer iPort.pg5)# map pool iPool4 type lossless reserved 7k xoff 2k xon 1k shared max 20`	
Related Commands		
Note	• Xon and Xoff values are in KB and valid only for “lossless” type • It is possible to use “K” and “M” quantifiers to set reserved size	

bind switch-priority

bind switch-priority <list-of-switch-priorities>

Bind a switch priority (SP) to an ingress buffer.
The no form of the command resets the values to their default.

Syntax Description	list-of-switch-priorities Possible values: 0-7
Default	According to system default OOB configuration
Configuration Mode	config interface ethernet ingress-buffer
History	3.6.1002
Role	admin
Example	`switch (config  1/1 ingress-buffer iPort.pg1)# bind switch-priority 0 1`
Related Commands	
Note	

description

description <description>

Configures buffer description.
The no form of the command resets the values to their default.

Syntax Description	description	Text string
Default	“”	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.1002	
Role	admin	
Example	`switch (config  1/1 ingress-buffer iPort.pg1)#` `description example`	
Related Commands		
Note		

pool mc-buffer

pool <pool-name> mc-buffer <buffer> reserved <reserved> shared <shared units> <shared-size>
no pool <pool-name> mc-buffer

Maps MC-buffer to specified egress pool and sets its reserved and shared sizes.
The no form of the command resets the values to their default.

Syntax Description	mc-buffer	Buffer can have the values mc.sp0, mc.sp1…mc.sp7
	reserved	The amount of shared memory for this buffer
	shared	The amount of shared memory for this buffer • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 … 1, 2, 4, … 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite
Default	N/A	
Configuration Mode	config config interface ethernet egress-buffer	
History	3.6.1002	
	3.6.5000	Added “size” parameter and note
Role	admin	
Example	`switch (config)# pool ePool4 mc-buffer mc.sp6 reserved 3k shared size 2K`	
Related Commands		

Note	• The qualifiers “K” and “M” may be used to set reserved and shared size
	• The units alpha, max, size is presented to the user according to the pool type “static”, “dynamic” and “size”:
	  • Alpha when pool type is dynamic and size is defined in bytes
	  • Max when pool type is static and size is defined in bytes
	  • Size when pool type is static and size is infinite	

pool description

pool <pool-name> description <description>
no pool <pool-name> description

Configures the buffer description of a specific pool-name.
The no form of the command resets the values to their default.

Syntax Description	pool-name	Possible values: • ePool0 … ePool7 • iPool0 … iPool7
	description	String text (20 character max)
Default	“”	
Configuration Mode	config	
History	3.6.1002	
	3.6.5000	
Role	admin	
Example	`switch (config)# pool iPool6 description mapped-to-pg3`	
Related Commands		
Note		

cable-length

cable-length [<meters>]

Configures the cable length in meters for the given port.

Syntax Description	meters	Cable length in meters Range: 5-100,000
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/4)# cable-length 10`	
Related Commands		
Note	• The user may use the quantifier “K” to indicate kilometers (e.g. “cable-length 5K”) • This command is used to calculate the required buffer to sustain the delay caused by the cable length	

show buffers mode

	show buffers mode Displays current mode for shared buffers.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin
Example	`switch (config)# show buffers mode` `Current mode: user mode`
Related Commands	
Note	

show buffers status

show buffers status [interfaces ethernet <slot>/<port>]

Displays buffer usage status.

Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated Example
	3.6.6000	Updated Example
Role	admin	

Example

```
switch (config)# show buffers status interfaces ethernet 1/9
-------------------------------------------------------------------------------
  Interface  Buffer           Pool       Resv      Shared        Usage     MaxUsage
                                        [Byte]    [%/a/Byte]    [Byte]    [Byte]
-------------------------------------------------------------------------------
  Eth1/9     iPort.iPool0     iPool0     10.0K     alpha 8       0         0
  Eth1/9     iPort.iPool1     iPool1     0         alpha 0       0         0
  Eth1/9     iPort.iPool2     iPool2     0         alpha 0       0         0
  Eth1/9     iPort.iPool3     iPool3     0         alpha 0       0         0
  Eth1/9     iPort.iPool4     iPool4     0         alpha 0       0         0
  Eth1/9     iPort.iPool5     iPool5     0         alpha 0       0         0
  Eth1/9     iPort.iPool6     iPool6     0         alpha 0       0         0
  Eth1/9     iPort.iPool7     iPool7     0         alpha 0       0         0
  Eth1/9     iPort.iPoolCtrl  iPoolCtrl  0         alpha 8       0         0
  Eth1/9     iPort.pg0        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg1        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg2        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg3        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg4        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg5        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg6        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg7        iPool0     0         alpha 8       0         0
  Eth1/9     iPort.pg9        iPoolCtrl  10.0K     alpha 8       0         0
  Eth1/9     ePort.ePool0     ePool0     5.0K      alpha 8       0         0
  Eth1/9     ePort.ePool1     ePool1     0         alpha 0       0         0
  Eth1/9     ePort.ePool2     ePool2     0         alpha 0       0         0
  Eth1/9     ePort.ePool3     ePool3     0         alpha 0       0         0
  Eth1/9     ePort.ePool4     ePool4     0         alpha 0       0         0
  Eth1/9     ePort.ePool5     ePool5     0         alpha 0       0         0
  Eth1/9     ePort.ePool6     ePool6     0         alpha 0       0         0
  Eth1/9     ePort.ePool7     ePool7     0         alpha 0       0         0
  Eth1/9     ePort.mc         ePool15    5.0K      90.0K         0         0
  Eth1/9     ePort.ePoolCtrl  ePoolCtrl  0         alpha 8       0         0
  Eth1/9     ePort.tc0        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc1        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc2        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc3        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc4        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc5        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc6        ePool0     1.0K      alpha 8       0         0
  Eth1/9     ePort.tc7        ePool0     1.0K      alpha 8       0         0
```

Related Commands

Note

show buffers details

show buffers details [<id>]

Displays buffer status in details.

Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated Example
	3.7.1000	Updated Example
Role	admin	

Example

```
switch (config)# show buffers details

Flags:
  Y: Lossy
  L: Lossless
  S: Static
  D: Dynamic

  Shared size is in percent/Bytes for static pool and in alphas for dynamic pool

Interface Eth1/1:
  ------------------------------------------------------------------------------------------
  Buffer              Resv      Xoff      Xon       Shared        Pool          Description
                      [Byte]    [Byte]    [Byte]    [%/a/Byte]
  ------------------------------------------------------------------------------------------
  iPort.iPool0(Y)     10.0K     -         -         alpha 8       iPool0(D)
  iPort.iPool1(Y)     0         -         -         alpha 0       iPool1(D)
  iPort.iPool2(Y)     0         -         -         alpha 0       iPool2(D)
  iPort.iPool3(Y)     0         -         -         alpha 0       iPool3(D)
  iPort.iPool4(Y)     0         -         -         alpha 0       iPool4(D)
  iPort.iPool5(Y)     0         -         -         alpha 0       iPool5(D)
  iPort.iPool6(Y)     0         -         -         alpha 0       iPool6(D)
  iPort.iPool7(Y)     0         -         -         alpha 0       iPool7(D)
  iPort.iPoolCtrl(Y)  0         -         -         alpha 8       iPoolCtrl(D)
  iPort.pg0(Y)        0         -         -         alpha 8       iPool0(D)
  iPort.pg1(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg2(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg3(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg4(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg5(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg6(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg7(Y)        0         -         -         alpha 0       iPool0(D)
  iPort.pg9(Y)        10.0K     -         -         alpha 8       iPoolCtrl(D)
  ePort.ePool0        10.0K     -         -         alpha 8       ePool0(D)
  ePort.ePool1        0         -         -         alpha 0       ePool1(D)
  ePort.ePool2        0         -         -         alpha 0       ePool2(D)
  ePort.ePool3        0         -         -         alpha 0       ePool3(D)
  ePort.ePool4        0         -         -         alpha 0       ePool4(D)
  ePort.ePool5        0         -         -         alpha 0       ePool5(D)
  ePort.ePool6        0         -         -         alpha 0       ePool6(D)
  ePort.ePool7        0         -         -         alpha 0       ePool7(D)
  ePort.mc            10.0K     -         -         90.0K         ePool15(S)
  ePort.ePoolCtrl     0         -         -         alpha 8       ePoolCtrl(D)
  ePort.tc0           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc1           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc2           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc3           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc4           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc5           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc6           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc7           1.0K      -         -         alpha 8       ePool0(D)
  ePort.tc16          1.0K      -         -         alpha 8       ePoolCtrl(D)
```

```
switch-priority to Buffers mapping:
   ----------------------------
   Switch-priority    Buffer
   ----------------------------
   0                  iPort.pg0
   1                  iPort.pg0
   2                  iPort.pg0
   3                  iPort.pg0
   4                  iPort.pg0
   5                  iPort.pg0
   6                  iPort.pg0
   7                  iPort.pg0
```

Related Commands

Note

show buffers pools

show buffers pools [pool-name]

Displays buffer pool statistics.

Syntax Description	pool-name	• iPool0-iPool7 • ePool0-ePool7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated example output
Role	admin	
Example		

```
switch (config)# show buffers pools
Flags: S - Static, D - Dynamic

------------------------------------------------------------------------
Pool       Direction     Size     Usage    MaxUsage     Description
                         [Byte]   [Byte]   [Byte]
------------------------------------------------------------------------
iPool0     ingress(D)    13.2M    0        576          Lossy-default
iPool1     ingress(D)    0        0        0
iPool2     ingress(D)    0        0        0
iPool3     ingress(D)    0        0        0
iPool4     ingress(D)    0        0        0
iPool5     ingress(D)    0        0        0
iPool6     ingress(D)    0        0        0
iPool7     ingress(D)    0        0        0
iPoolCtrl  ingress(D)    256.0K   0        0            Control
ePool0     egress(D)     13.2M    0        0            Default
ePool1     egress(D)     0        0        0
ePool2     egress(D)     0        0        0
ePool3     egress(D)     10.0K    0        0
ePool4     egress(D)     0        0        0
ePool5     egress(D)     0        0        0
ePool6     egress(D)     0        0        0
ePool7     egress(D)     0        0        0
ePool15    egress(S)     inf      0        0            Multicast
ePoolCtrl  egress(D)     256.0K   0        0            Control
```

Related Commands	
Note	When advanced buffer management is disabled, the “Description” field specifies the e/iPool’s relevant traffic pool name.

show buffers pools mc-buffers

show buffers pools [<pool-name>] mc-buffers

Displays global multicast buffers usage status.

Syntax Description	pool-name	Possible values: ePool0 … ePool7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	
Example	`switch (config)# show buffers pools ePool4 mc-buffers`	

```
------------------------------------------------------------------------
  MC-Buffer   Pool        Resv      Shared          Usage       MaxUsage
                          [Byte]    [%/a/Byte]      [Byte]      [Byte]
------------------------------------------------------------------------
  mc.sp0      ePool0      0         alpha 1/4       0           0
  mc.sp1      ePool0      0         alpha 1/4       0           0
  mc.sp2      ePool0      0         alpha 1/4       0           0
  mc.sp3      ePool0      0         alpha 1/4       0           0
  mc.sp4      ePool0      0         alpha 1/4       0           0
  mc.sp5      ePool0      0         alpha 1/4       0           0
  mc.sp6      ePool0      0         alpha 1/4       0           0
  mc.sp7      ePool0      0         alpha 1/4       0           0
```

Related Commands	
Note	

show traffic pool

show traffic pool [<name>]

Displays state and configuration information for a given traffic pool.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin

Example	`switch (config)# show traffic pool`
	
	`---------------------------------------------------------------------------------------------`
	`Traffic Pool  Type  Memory [%]  Switch Priorities  Memory actual [Bytes]  Usage [KB]  Max Usage [Bytes]`
	`---------------------------------------------------------------------------------------------`
	`lossless-default (RO)  lossless  auto  0  0  0`
	`lossy-default  lossy  auto  0, 1, 2, 3, 4, 5, 6, 7  13.7M  0  0`
	
	`Exception list:`
	`N/A`

Related Commands	
Note	• Omission of traffic pool name displays information about all existing traffic pools • The “Exception list” section displays messages to indicate unrecommended configuration. See Section 5.17.3.8, “Exceptions to Legal Shared Buffer Configuration,” on page 1112 for more details.

show traffic pool

show traffic pool <name> <device/port>

Displays state and configuration information for the buffers on a given port related to a given traffic pool.

Syntax Description	<device/port> Ethernet interface
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Role	admin
Example	

```
switch (config)# show  traffic pool lossy-default  1/1

-----------------------------------------------------
Switch-priority   Ingress buffer    Egress buffer
-----------------------------------------------------
0                 iPort.pg0         ePort.tc0
1                 iPort.pg0         ePort.tc1
2                 iPort.pg0         ePort.tc2
3                 iPort.pg0         ePort.tc3
4                 iPort.pg0         ePort.tc4
5                 iPort.pg0         ePort.tc5
6                 iPort.pg0         ePort.tc6
7                 iPort.pg0         ePort.tc7

-------------------------------------------------------------------------
Name                  Memory percent    Size (bytes)      Usage (bytes)
-------------------------------------------------------------------------
lossy-default         auto              13.7M              0

-------------------------------------------------------------------
Ingress buffer  Headroom size (bytes)  Xon (bytes)  Xoff (bytes)
-------------------------------------------------------------------
iPort.pg0       N/A                     N/A          N/A

-----------------------------------------------------------
Direction    Pool Usage (bytes)   Pool Max Usage (bytes)
-----------------------------------------------------------
Ingress      0                    0
Egress       0                    0

Exception list:
N/A
```

Related Commands	
Note	The “Exception list” section displays messages to indicate unrecommended configuration. See Section 5.17.3.8, “Exceptions to Legal Shared Buffer Configuration,” on page 1112 for more details.

5.18 Storm Control

Storm control may be enabled on L2 Ethernet ports, LAGs, and MLAGs to monitor inbound traffic to prevent disruptions caused by a broadcast, multicast, or unicast traffic storm on the physical interfaces.

Storm control utilizes a bandwidth-based method to measure traffic where packets exceeding the percentage level specified by the user are dropped.

Users are able to monitor broadcast, unknown unicast, and unregistered multicast traffic while supporting differing thresholds for each type or monitor a summary of all the previously mentioned traffic with one threshold.

5.18.1 Commands

storm-control

storm-control {<broadcast | unreg-multicast | unknown-unicast> | all} {level <level> | { bits <bits> | bytes <bytes> | packets <packets> [k|m|g]}} [force]
[no] storm-control {<broadcast | unreg-multicast | unknown-unicast> | all}

The command enables Storm Control on selected interface.
The no form of the command disables Storm Control on selected interface.

Syntax Description	{broadcast|unreg-multi-cast|unknown-unicast|all}	• Each port can support broadcast, unregistered-multi-cast, unknown-unicast or all configurations. • All means one threshold level for all traffic types. It is not identical to configuring broadcast, unregis-tered-multicast and unknown-unicast together.
	{level <level> | { bits <bits> | bytes <bytes> | packets <packets> [k|m|g]}}	Storm control per traffic type may be configured with different thresholds. • Level – specifies threshold value in percentages from interface speed. • Bits – specifies threshold value in bits per second. Must be specified with multiplier k, m, or g. Avail-able ranges: [1k...999k][1m...999m][1g...200g]. • Bytes – specifies threshold value in bytes per second. May be specified with multiplier k, m, or g. Available ranges: [128...999][1k...999k][1m...999m][1g...25g]. • Packets – specifies threshold value in packets per sec-ond. May be specified with multiplier k, m, or g. Available ranges: [1...999][1k...999k][1m...999m][1g...2g].
Default	no storm control	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.4006	
	3.6.4110	Updated command syntax, default and configuration mode
	3.6.6000	Added “config interface mlag port channel” configura-tion mode
	3.7.00xx	Added bits/bytes/packets threshold types
Role	admin	

Example	`interface ethernet 1/1 storm-control broadcast bits 100 m`
	`interface ethernet 1/1 storm-control unknown-unicast level 50`
	`interface ethernet 1/1 storm-control unreg-multicast packets 900`
	`interface ethernet 1/2 storm-control all bytes 1 g`	

Related Commands	
Note	• "all" and other configurations are mutually exclusive • User can use the "force" parameter to resolve collision and apply new configuration • Storm Control can be configured on LAG but cannot be configured on LAG members • Storm Control cannot be configured on router ports • Storm Control cannot be configured on a port which takes part in monitoring sessions as destination port • Units are in 10^n. k = 1000 and not 1024.

show storm-control

show storm-control [<interface>]

The command displays the configuration levels and dropped packets for each traffic type.

Syntax Description	<Interface>	• Displays configuration and dropped packets on specified interface. • If interface is not specified, displays configuration and dropped packets on all interfaces.
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.6.4110	Updated example output.
	3.7.1000	Updated example output.
Role	admin	
Example	`r-qa-sw-eth-85 [standalone: master] (config) # show  storm-control`	
Related Commands		
Note		

```
Interface Eth1/8:
  Broadcast                        : 10%
  Broadcast packets dropped        : 0
  Unreg-Mcast                      : N/A
  Unreg-Mcast packets dropped      : N/A
  Unkn-Ucast                       : N/A
  Unkn-Ucast packets dropped       : N/A
  All traffic types                : N/A
  All traffic types packets dropped: N/A
```

5.19 Head-of-Queue Lifetime Limit

Head-of-queue (HoQ) lifetime limit (HLL) is a mechanism which allows discarding packets attempting to be transmitted after HLL time from the time that they were ready to be transmitted at the head of the scheduling group.

When HLL_packet2Stall (7 as default) packets encounter HLL drop, the scheduling group enters a stall state. During that state all packets to the sub-group are discarded. The subgroup exits stall state after HLL_time*8.

A counter called HoQ discard packets counts the number of discarded packets due to HLL.

5.19.1 Commands

hll

hll <max-time>
no hll

Configures HLL time on this interface.
The no form of the command resets HLL time to its default value.

Syntax Description	max-time	Possible values: • <4 | 16 | 32 | 64 | 128 | 256 | 512>ms • <1 | 2>sec • "inf" to disable HLL
Default	512ms	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Role	admin	
Example	`switch (config  1/10)# hll 512ms`	
Related Commands		
Note		

5.20 User Defined Keys

User defined keys (UDKs) allow defining custom byte keys—that is, groups of bytes that can be matched to a predefined point in the packet (an extraction point, e.g. the start of a MAC header, or an IP header)—which is useful when wanting to make a match with a part of the packet which does not have a dedicated key.

The maximum number of UDKs is 4.

An extraction point may be defined for each packet type in a UDK. For each extraction point, an offset (from the beginning of the extraction) is defined.

To be able to modify a UDK after attaching it to an ACL rule, it is first necessary to un-match the UDK from the ACL, and then change the match mode of the UDK to none using the command “no udk match mode”.

Defining a UDK affects the throughput for packets equal or smaller than 128 bytes.

5.20.1 Configuring UDK

➢ ***To set UDK with ACL on a specific field:***

Step 1. Define new user defined key called ipv4_udk. Run:

```
switch (config) # udk ipv4_udk
switch (config) # exit
```

Step 2. Set user defined key ipv4_udk to match on IPV4 header in offset 4 bytes from start of header. Run:

```
switch (config) # udk ipv4_udk extraction point mode l3 packet type ipv4 extraction
point start-of-header offset 4
```

Step 3. Set the len (in bytes) of the field to match on. Run:

```
switch (config) # udk ipv4_udk len 2
```

Step 4. Set the user defined key to work with access list. Run:

```
switch (config) # udk ipv4_udk match mode acl
```

Step 5. Define new access list table called my_acl_table. Run:

```
switch (config) # ipv4-udk access-list my_acl_table
switch (config) # exit
```

Step 6. Set new rule on the access list table with the previously defined user defined key to match 0x1234. Run:

```
switch (config) # ipv4-udk access-list my_acl_table permit ip any any udk ipv4_udk
0x1234
```

Step 7. Bind the access list table to an ethernet interface. Run:

```
switch (config) # interface ethernet 1/1 ipv4-udk port access-group my_acl_table
```

5.20.2 Commands

udk

udk <udk-name>
no udk <udk-name>

Creates user defined key.
The no form of the command deletes user defined key.

Syntax Description	udk-name	String
Default	N/A	
Configuration Mode	config	
History	3.6.5000	
Role	admin	
Example	`switch (config)# udk udk_name` `switch (config udk udk_name)#`	
Related Commands		
Note	Defining UDK affects the throughput for packets equal or smaller than 128 bytes.	

match mode

match mode <match-mode>
no match mode

Configures user defined key match mode.
The no form of the command resets this parameter to its default.

Syntax Description	match-mode	Possible values: • acl • all
Default	None	
Configuration Mode	config udk	
History	3.6.5000	
Role	admin	
Example	`switch (config udk udk_name)# match mode all`	
Related Commands	udk <udk-name>	
Note		

extraction point

extraction point mode <mode> [packet type <type> [extraction point <point> [offset <offset>]]]

Configures user-defined key extraction point mode.

Syntax Description	mode	Possible values: • l2 • l3 • l4
	packet type	Sets user defined key packet type. Possible values: • For L2: l2 • For L3: arp; ipv4; ipv6 • For L4: udp
	extraction point	Sets user defined key extraction point. Possible values for: • l2: l2-ether-type; start-of-header • arp: start-of-header • ipv4; ipv6: start-of-header; start-of-payload • udp: start-of-payload
	offset	Sets user defined key extraction point offset Range: 0-126 (even values)

Default	• Mode: l3
	• Default extraction point per packet type:
	  • L2: start-of-header
	  • ARP; IPv4; IPv6: start-of-header
	  • UDP: start-of-payload
	• Offset: 0	

Configuration Mode	config udk	
History	3.6.5000	
Role	admin	
Example	`switch (config udk udk_name)# extraction point mode l3 packet type ipv4 extraction point start-of-header offset 2`	
Related Commands	udk <udk-name>	
Note		

len

len <length>

Configures user-defined key length.

Syntax Description	length	Range: 1-4
Default	4	
Configuration Mode	config udk	
History	3.6.5000	
Role	admin	
Example	`switch (config udk udk_name)# len 4`	
Related Commands	udk <udk-name>	
Note		

show udk

show udk [<udk-name>]

Displays summary for user-defined keys.

Syntax Description	udk-name	Displays information about specific UDK
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Role	admin	

Example	`switch (config)# show udk`
	
	`UDK name: udk_name`
	`Match mode: none`
	`Length: 4`
	`Extraction mode: l3`
	`IPv4 extraction point: start-of-header`
	`IPv4 offset: 22`
	`IPv6 extraction point: start-of-header`
	`IPv6 offset: 0`
	`ARP extraction point: start-of-header`
	`ARP offset: 0`	

Related Commands	udk <udk-name>	
Note		

5.21 Port Mirroring

Port mirroring enables data plane monitoring functionality which allows the user to send an entire traffic stream for testing. Port mirroring sends a copy of packets of a port's traffic stream, called "mirrored port", into an analyzer port. Port mirroring is used for network monitoring. It can be used for intrusion detection, security breaches, latency analysis, capacity and performance matters, and protocol analysis.

Figure 27 provides an overview of the mirroring functionality.

Figure 27: Overview of Mirroring Functionality

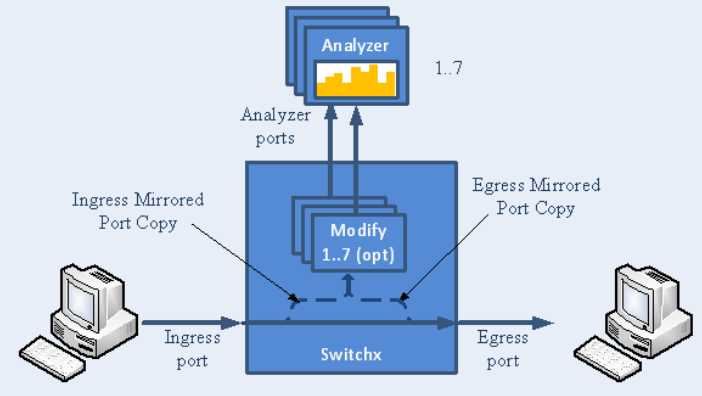

There is no limitation on the number of mirroring sources and more than a single source can be mapped to a single analyzer destination.

5.21.1 Mirroring Sessions

Port mirroring is performed by configuring mirroring sessions. A session is an association of a mirror port (or more) and an analyzer port.

Figure 28: Mirror to Analyzer Mapping

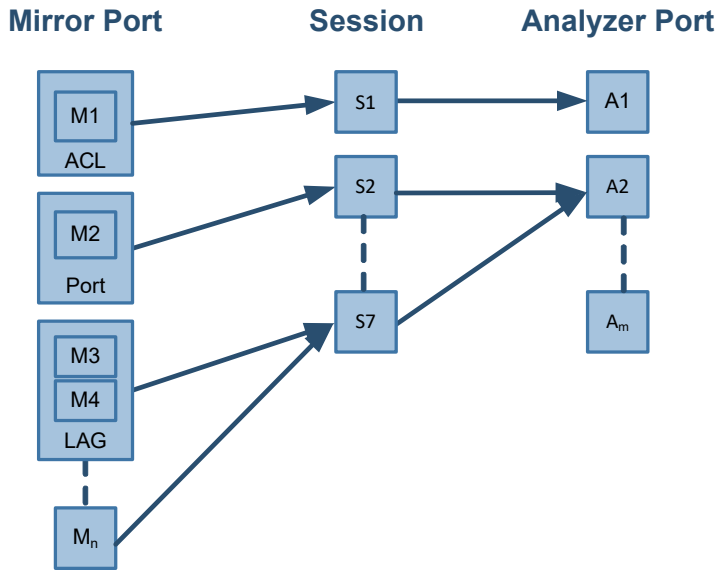

A mirroring session is a monitoring configuration mode that has the following parameters:

Table 59 - Mirroring Parameters

Parameter	Description	Access
Source interface(s)	List of source interfaces to be mirrored.	RW
Destination interface	A single analyzer port through which all mirrored traffic egress.	RW
Header format	The format and encapsulation of the mirrored traffic when sent to analyzer.	RW
Truncation	Enabling truncation segments each mirrored packet to 64 bytes.	RW
Congestion control	Controls the behavior of the source port when destination port is congested.	RW
Admin state	Administrative state of the monitoring session.	RW

5.21.1.1 Source Interface

The source interface (mirror port) refers to the interface from which the traffic is monitored. Port mirroring does not affect the switching of the original traffic. The traffic is simply duplicated and sent to the analyzer port. Traffic in any direction (either ingress, egress or both) can be mirrored.

There is no limitation on the number of the source interfaces mapped to a mirroring session.

Ingress and egress traffic flows of a specific source interface can be mapped to two different sessions.

LAG

The source interface can be a physical interface or a LAG.

Port mirroring can be configured on a LAG interface but not on a LAG member. When a port is added to a mirrored LAG it inherits the LAG's mirror configuration. However, if port mirroring configuration is set on a port, that configuration must be removed prior to adding the port to a LAG interface.

When a port is removed from a LAG, the mirror property is switched off for that port.

Control Protocols

All control protocols captured on the mirror port are forwarded to the analyzer port in addition to their normal treatment. For example LACP, STP, and LLDP are forwarded to the analyzer port in addition to their normal treatment by the CPU.

Exceptions to the behavior above are the packets that are being handled by the MAC layer, such as pause frames.

5.21.1.2 Destination Interface

The destination interface is an analyzer port to which mirrored traffic is directed. The mirrored packets are duplicated, optionally modified, and sent to the analyzer port. Spectrum™ platforms support up to only 2 analyzer ports, where any mirror port can be mapped to any analyzer port and more than a single mirror port can be mapped to a single analyzer port.

Packets can be forwarded to any destination using the command `destination interface`.

The analyzer port supports status and statistics as any other port.

LAG

The destination interface cannot be a member of LAG when the header format is local.

Control Protocols

The destination interface may also operate in part as a standard port, receiving and sending out non-mirrored traffic. When the header format is configured as a local port, ingress control protocol packets that are received by the local analyzer port get discarded.

Advanced MTU Considerations

The analyzer port, like its counterparts, is subject to MTU configuration. It does not send packets longer than configured.

When the analyzer port sends encapsulated traffic, the analyzer traffic has additional headers and therefore longer frame. The MTU must be configured to support the additional length, otherwise, the packet is truncated to the configured MTU.

The system on the receiving end of the analyzer port must be set to handle the egress traffic. If it is not, it might discard it and indicate this in its statistics (packet too long).

5.21.1.3 Header Format

Ingress traffic from the source interface can be manipulated in several ways depending on the network layout using the command `header-format`.

If the analyzer system is directly connected to the destination interface, then the only parameters that can be configured on the port are the MTU, speed and port based flow control. Priority flow control is not supported is this case. However, if the analyzer system is indirectly connected to the destination interface, there are two options for switching the mirrored data to the analyzer system:

- A VLAN tag may be added to the Ethernet header of the mirrored traffic
- An Ethernet header can be added with include a new destination address and VLAN tag

It must be taken into account that adding headers increases packet size.

Figure 29: Header Format Options

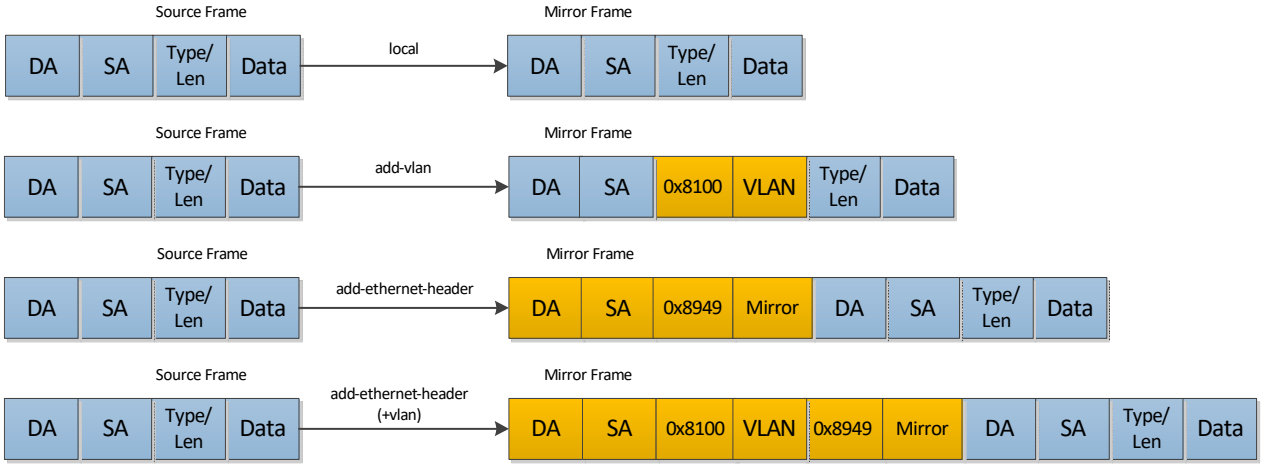

5.21.1.4 Congestion Control

The destination ports might receive pause frames that lead to congestion in the switch port. In addition, too much traffic directed to the analyzer port (for example 40GbE mirror port is directed into 10GbE analyzer port) might also lead to congestion.

In case of congestion:

- When best effort mode is enabled on the analyzer port, Spectrum drops excessive traffic headed to the analyzer port using tail drop mechanism, however, the regular data (mirrored data heading to its original port) does not suffer from a delay or drops due to the analyzer port congestion.
- When the best effort mode on the analyzer port is disabled, the Spectrum does not drop the excessive traffic. This might lead to buffer exhaustion and data path packet loss.

The default behavior in congestion situations is to drop any excessive frames that may clog the system.

ETS, PFC and FC configurations do not apply to the destination port.

5.21.1.5 Truncation

When enabled, the system can truncate the mirrored packets into smaller 64-byte packets (default) which is enough to capture the packets' L2 and L3 headers.

5.21.2 Configuring Mirroring Sessions

Figure 30 presents two network scenarios with direct and remote connectivity to the analyzer equipment. Direct connectivity is when the analyzer is connected to the analyzer port of the switch. In this case there is no need for adding an L2 header to the mirrored traffic. Remote connectivity is when the analyzer is indirectly connected to the analyzer port of the switch. In this situation, adding an L2 header may be necessary depending on the network's setup.

Figure 30: Mirroring Session

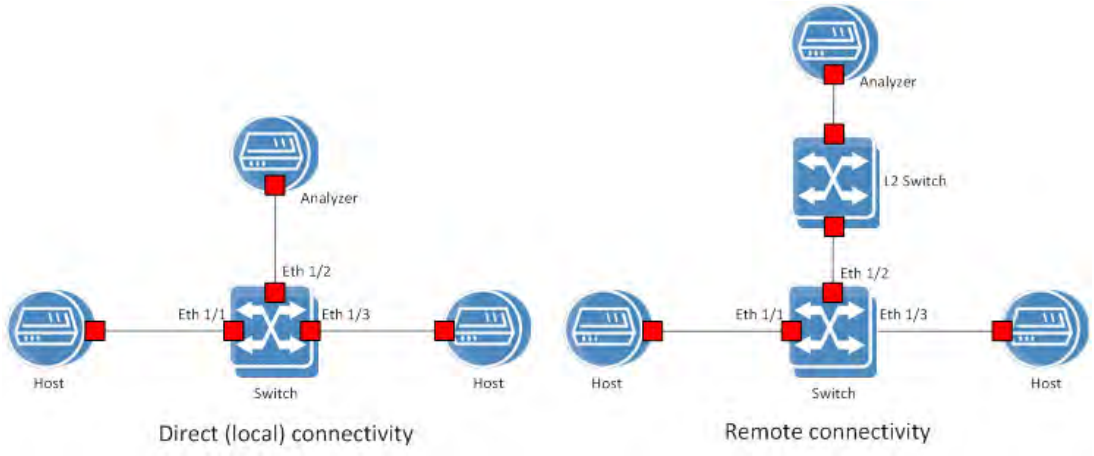

➢ ***To configure a mirroring session:***

Step 1. Create a session. Run:

```
switch (config) # monitor session 1
```

This command enters a monitor session configuration mode. Upon first implementation the command also creates the session.

Step 2. Add source interface(s). Run:

```
switch (config monitor session 1) # add source  1/1 direction both
```

Step 3. Add destination interface. Run:

```
switch (config monitor session 1) # destination  1/2
```

Step 4. (Optional) Set header format. Run:

```
switch (config monitor session 1) # header-format add-ethernet-header destination-mac
00:0d:ec:f1:a9:c8 add-vlan 10 priority 5 traffic-class 2
```

For remote connectivity use the header formats `add-vlan` or `add-ethernet-header`. For local connectivity, use `local`.

Step 5. (Optional) Truncate the mirrored traffic to 64-byte packets. Run:

```
switch (config monitor session 1) # truncate
```

Step 6. (Optional) Set congestion control. Run:

```
switch (config monitor session 1) # congestion pause-excessive-frames
```

The default for this command is to drop excessive frames. The `pause-excessive-frames` option uses flow control to regulate the traffic from the source interfaces.

If the option `pause-excessive-frame` is selected, make sure that flow control is enabled on **all** source interfaces on the ingress direction of the monitoring session using the command `flowcontrol` in the interface configuration mode.

Step 7. Enable the session. Run:

```
switch (config monitor session 1) # no shutdown
```

5.21.3 Verifying Mirroring Sessions

➢ ***To verify the attributes of a specific mirroring session:***

```
switch (config) # show monitor session 1
Session 1:
  Admin:  Enable
  Status: Up
  Truncate: Enable
  Destination interface: eth1/2
  Congestion type: pause-excessive-frames
  Header format: add-ethernet-header
  -switch priority: 5

Source interfaces
--------------------
Interface  Direction
--------------------
eth1/1     both
```

➢ ***To verify the attributes of running mirroring sessions:***

```
switch (config) # show monitor session summary
Flags: i ingress, e egress, b both

------------------------------------------------------------
Session  Admin      Status  Mode       Destination  Source
------------------------------------------------------------
1        Enable     Up      add-eth    eth1/2       eth1/1(b)
2        Disable    Down    add-vlan   eth1/2       eth1/8(i), po1(e)
3        Enable     Up      add-eth    eth1/5       eth1/18(e)
7        Disable    Down    local
```

5.21.4 Commands

5.21.4.1 Config

monitor session

monitor session <session-id>
no monitor session <session-id>

Creates session and enters monitor session configuration mode upon using this command for the first time.
The no form of the command deletes the session.

Syntax Description	session-id	The monitor session ID Range is: 1-2
Default	N/A	
Configuration Mode	config	
History	3.3.3500	
Role	admin	
Example	`switch (config)# monitor session 1` `switch (config monitor session 1)#`	
Related Commands		
Note		

5.21.4.2 Config Monitor Session

destination interface

destination interface <type> <number> [force]
no destination interface

Sets the egress interface number.
The no form of the command deletes the destination interface.

Syntax Description	interface	Sets the interface type and number (e.g. ethernet 1/2)
	force	Eliminates the need to shutdown the port prior to the operation
Default	no destination interface	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.3.4100	Added force argument
	3.6.4006	Added note
Role	admin	
Example	`switch (config monitor session 1) # destination  1/2`	
Related Commands		
Note	• Port cannot be used as destination port in monitor session when storm-control is configured on port. • Force command cannot remove storm-control configuration. Error output: `Configuration error, storm control is configured on port` • When removing an interface from a monitor session it gains the default attributes of Ethernet ports	

shutdown

	shutdown **no shutdown** Disables the session. The no form of the command enables the session.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config monitor session
History	3.3.3500
Role	admin
Example	`switch (config monitor session 1) # no shutdown` `switch (config monitor session 1)#`
Related Commands	
Note	

add source interface direction

	add source interface <type> <number> direction <d-type> **no source interface <type> <number>** Adds a source interface to the mirrored session. The no form of the command deletes the source interface.	
Syntax Description	interface <type> <number>	Interface type to configure
	direction <d-type>	Configures the direction of the mirrored traffic. The options are as follows: • egress – monitors egress traffic • ingress – monitors ingress traffic • both – monitors egress and ingress traffic
Default	N/A	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.5.1000	Updated
Role	admin	
Example	`switch (config monitor session 1) # add source  1/1 direction ingress`	
Related Commands		
Note	• If mirroring is configured in one direction (e.g. ingress) on an interface and then is configured in the other direction (e.g. egress), then the ultimate setting is “both” • Only ingress traffic mirroring is supported	

header-format

header-format {local [traffic-class <tc>] | add-vlan <vlan-id> [priority <prio>] [switch-priority <sp>] | add-ethernet-header destination-mac <mac-address> [add-vlan <vlan-id> [priority <prio>]] [traffic-class <tc>]}
no header-format

Sets the header format of the mirrored traffic.
The no form of the command resets the parameter values back to default.

Syntax Description	local	The mirrored header of the frame is not changed.
	switch-priority <sp>	Changes the egress switch priority of the frame. Range: 0-15.
	add-vlan <vlan-id>	An 802.1q VLAN tag is added to the frame.
	priority <prio>	The priority to be added to the Ethernet header. Range: 0-7.
	add-ethernet-header	Adds an Ethernet header to the mirrored frame.
	destination-mac	The destination MAC address of the added Ethernet frame.
Default	no-change vlan 1 priority 0 traffic-class 0	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.5.1000	Added switch-priority parameter
Role	admin	
Example	`switch (config monitor session 1) # header-format add-ethernet-header destination-mac 00:0d:ec:f1:a9:c8 add-vlan 10 priority 5 traffic-class 2`	
Related Commands		
Note	If add-ethernet-header is used, the source MAC address is the one attached to the switch	

truncate

	truncate **no truncate** Truncates the mirrored frames to 64-byte packets. The no form of the command disables truncation.
Syntax Description	N/A
Default	no truncate
Configuration Mode	config monitor session
History	3.3.3500
Role	admin
Example	`switch (config monitor session 1) # truncate` `switch (config monitor session 1)#`
Related Commands	
Note	This command applies for all sessions on the same analyzer port.

congestion

congestion [drop-excessive-frames | pause-excessive-frames]
no congestion

Sets the system's behavior when congested
The no form of the command disables truncation.

Syntax Description	drop-excessive-frames	Drops excessive frames.
	pause-excessive-frames	Pauses excessive frames.
Default	drop-excessive-frames	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.3.4000	Added Syntax Description.
Role	admin	
Example	`switch (config monitor session 1) # congestion pause-excessive-frames` `switch (config monitor session 1)#`	
Related Commands		
Note	This command applies for all sessions on the same analyzer port.	

5.21.4.3 Show

show monitor session

show monitor session <session-id>

Displays monitor session configuration and status.

Syntax Description	session-id	The monitor session ID Range: 1-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.5.1000	Updated Note section
	3.6.5000	Updated Example
Role	admin	

Example	`switch (config) # show monitor session 1`
	`Session 1:`
	`  Admin:  Disable`
	`  Status: Down`
	`  Truncate:   Disable`
	`  Destination interface: N/A`
	`  Congestion type: drop-excessive-frames`
	`  Header format: local`
	`  -switch priority: 0`
	
	`Source interfaces`
	`--------------------`
	`Interface  Direction`
	`--------------------`
	`eth1/1     both`
	`switch (config) #`	

Related Commands		
Note		

show monitor session summary

	show monitor session summary Displays monitor session configuration and status summary.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.3500 3.6.5000 Updated Example
Role	admin

Example	`switch (config) # show monitor session summary`
	`Flags: i ingress, e egress, b both`
	
	`------------------------------------------------------------`
	`Session  Admin      Status  Mode       Destination  Source`
	`------------------------------------------------------------`
	`1        Disable    Down    local      N/A          eth1/1(b)`
	`2        Disable    Down    add-vlan   eth1/2       eth1/8(i)`
	`switch (config) #`

Related Commands	
Note	

5.22 sFlow

sFlow (ver. 5) is a procedure for statistical monitoring of traffic in networks. Onyx supports an sFlow sampling mechanism (agent), which includes collecting traffic samples and data from counters. The sFlow datagrams are then sent to a central collector.

The sampling mechanism must ensure that any packet going into the system has an equal chance of being sampled, irrespective of the flow to which it belongs. The sampling mechanism provides the collector with periodical information on the amount (and load) of traffic per interface by loading the counter samples into sFlow datagrams.

The sFlow packets are encapsulated and sent in UDP over IP. The UDP port number that is used is the standard 6343 by default.

Figure 31: sFlow Functionality Overview

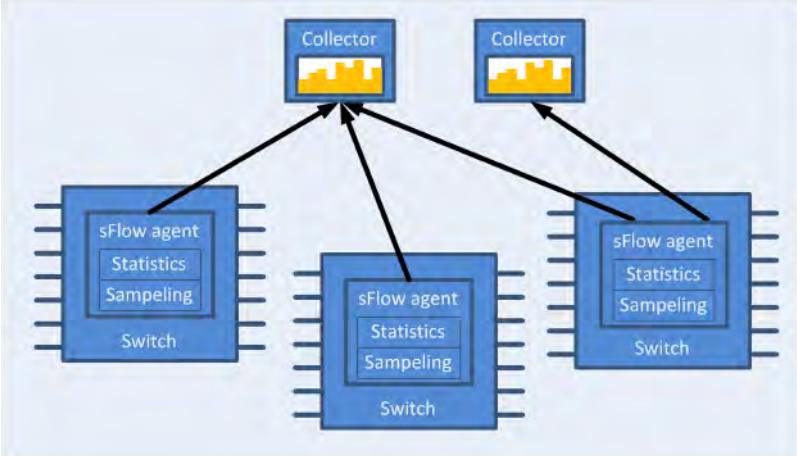

5.22.1 Flow Samples

The sFlow agent samples the data path based on packets.

Truncation and sampling rate are the two parameters that influence the flow samples. In case of congestion the flow samples can be truncated to a predefined size before it is assigned to the CPU. The truncation can be set to any value between 64 to 256 bytes with the default being 128 bytes.

The sampling rate can be adjusted by setting an average rate. The system assures that a random number of packets is sampled, however, the sample rate on average converges to the configured rate. Valid values range between 4000 to 16777215 packets.

5.22.2 Statistical Samples

The sFlow agent samples interface counters time based. Polling interval is configurable to any value between 5-3600 seconds with the default being 20 seconds.

The following statistics are gathered by the CPU:

Table 60 - List of Statistical Counters

Counter	Description
Total packets	The number of packets that pass through sFlow-enabled ports.
Number of flow samples	The number of packets that are captured by the sampling mechanism.
Number of statistic samples	The number of statistical samples.
Number of discarded samples	The number of samples that were discarded.
Number of datagrams	The number of datagrams that were sent to the collector.

5.22.3 sFlow Datagrams

The sFlow datagrams contain flow samples and statistical samples.

The sFlow mechanism uses IP protocol, therefore if the packet length is more than the interface MTU, it becomes fragmented by the IP stack. The MTU may also be set manually to anything in the range of 200-9216 bytes. The default is 1400 bytes.

5.22.4 Sampled Interfaces

sFlow must be enabled on physical or LAG interfaces that require sampling. When adding a port to a LAG, sFlow must be disabled on the port. If a port with enabled sFlow is configured to be added to a LAG, the configuration is rejected. Removing a port from a LAG disables sFlow on the port regardless of the LAG's sFlow status.

5.22.5 Configuring sFlow

➢ ***To configure the sFlow agent:***

Step 1. Unlock the sFlow commands. Run:

```
switch (config) # protocol sflow
```

Step 2. Enable sFlow on the system. Run:

```
switch (config) # sflow enable
```

Step 3. Enter sFlow configuration mode. Run:

```
switch (config) # sflow
switch (config sflow) #
```

Step 4. Set the central collector's IP. Run:

```
switch (config sflow) # collector-ip 10.10.10.10
```

Step 5. Set the agent-ip used in the sFlow header. Run:

```
switch (config sflow) # agent-ip 20.20.20.20
```

Step 6. (Optional) Set the sampling rate of the mechanism. Run:

```
switch (config sflow) # sampling-rate 16000
```

This means that one every 16000 packet gets collected for sampling.

Step 7. (Optional) Set the maximum size of the data path sample. Run:

```
switch (config sflow) # max-sample-size 156
```

Step 8. (Optional) Set the frequency in which counters are polled. Run:

```
switch (config sflow) # counter-poll-interval 19
```

Step 9. (Optional) Set the maximum size of the datagrams sent to the central collector. Run:

```
switch (config sflow) # max-datagram-size 1500
```

Step 10. Enable the sFlow agent on the desired interfaces. Run:

```
switch (config  1/1)# sflow enable
switch (config interface port-channel 1)# sflow enable
```

5.22.6 Verifying sFlow

➢ ***To verify the attributes of the sFlow agent:***

```
switch (config)# show sflow
sflow protocol: enabled
sflow: enabled
sampling-rate: 16000
max-sampled-size: 156
counter-poll-interval: 19
max-datagram-size: 1500
collector-ip: 10.10.10.10
collector-port: 6343
agent-ip: 20.20.20.20

ingress ports:
Interfaces:
Ethernet: eth1/1
Port-channel: po1

Statistics:
Total Samples: 2000
Number of flow samples: 1200
Estimated Number of flow discarded: 0
Number of statistic samples: 800
Number of datagrams: 300
```

5.22.7 Commands

5.22.7.1 Config

protocol sflow

	protocol sflow **no protocol sflow** Unhides the sFlow commands. The no form of the command deletes sFlow configuration and hides the sFlow commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.3500
Role	admin
Example	`switch (config) # protocol sflow` `switch (config) #`
Related Commands	
Note	

sflow enable (global)

	sflow enable **no sflow enable** Enables sFlow in the system. The no form of the command disables sFlow without deleting the configuration.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.3500
Role	admin
Example	`switch (config) # sflow enable` `switch (config) #`
Related Commands	
Note	

sflow

	sflow Enters sFlow configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.3500
Role	admin
Example	`switch (config) # sflow` `switch (config sflow) #`
Related Commands	
Note	

5.22.7.2 Config sFlow

sampling-rate

sampling-rate <rate>
no sampling-rate

Sets sFlow sampling ratio.
The no form of the command resets this parameter to its default value.

Syntax Description	rate	Sets the number of packets passed before selecting one for sampling. The range is 4000-16777215. Zero disables sampling.
Default	16000	
Configuration Mode	config sflow	
History	3.3.3500	
Role	admin	
Example	`switch (config sflow) # sampling-rate 16111` `switch (config sflow) #`	
Related Commands		
Note		

max-sample-size

max-sample-size <packet-size>
no max-sample-size

Sets the maximum size of sampled packets by sFlow.
The no form of the command resets the parameter to its default value.

Syntax Description	packet-size	The sampled packet size. The range is 64-256 bytes.
Default	128 bytes	
Configuration Mode	config sflow	
History	3.3.3500	
Role	admin	
Example	`switch (config sflow) # max-sample-size 165` `switch (config sflow) #`	
Related Commands		
Note	Sampled payload beyond the configured size is discarded.	

counter-poll-interval

counter-poll-interval <seconds>
no counter-poll-interval

Sets the sFlow statistics polling interval.
The no form of the command resets the parameter to its default value.

Syntax Description	seconds	The sFlow statistics polling interval in seconds. Range is 5-3600 seconds. Zero disables the statistic polling.
Default	20 seconds	
Configuration Mode	config sflow	
History	3.3.3500	
Role	admin	
Example	`switch (config sflow) # counter-poll-interval 30` `switch (config sflow) #`	
Related Commands		
Note		

max-datagram-size

max-datagram-size <packet-size>
no max-datagram-size

Sets the maximum sFlow packet size to be sent to the collector.
The no form of the command resets the parameter to its default value.

Syntax Description	packet-size	The packet size of the packet being sent to the collector. The range is 200-9216 bytes.
Default	1400 bytes	
Configuration Mode	config sflow	
History	3.3.3500	
Role	admin	
Example	`switch (config sflow) # max-datagram-size 9216` `switch (config sflow) #`	
Related Commands		
Note	This packet contains the data sample as well as the statistical counter data.	

collector-ip

collector-ip <ip-address> [udp-port <udp-port-number>]
no collector-ip [<ip-address> udp-port]

Sets the collector's IP.
The no form of the command resets the parameters to their default values.

Syntax Description	ip-address	The collector IP address.
	udp-port <udp-port-number>	Sets the collector UDP port number.
Default	ip-address: 0.0.0.0 udf-port-number: 6343	
Configuration Mode	config sflow	
History	3.3.3500	
Role	admin	
Example	`switch (config sflow) # collector-ip 10.10.10.10` `switch (config sflow) #`	
Related Commands		
Note		

agent-ip

agent-ip {<ip-address> | interface [ethernet <slot/port> | port-channel <channel-group>] | <if-name> | loopback <number> | vlan <id>}
no agent-ip

Sets the IP address associated with this agent.
The no form of the command resets the parameters to their default values.

Syntax Description	interface	Configures a specific ethernet/port-channel interface's agent IP.
	if-name	Interface name (e.g. mgmt0, mgmt1).
	ip-address	The sFlow agent's IP address (i.e. the source IP of the packet).
	loopback <number>	Loopback interface number. Range: 1-32.
	vlan <id>	Interface VLAN. Range: 1-4094.
Default	ip-address: 0.0.0.0	
Configuration Mode	config sflow	
History	3.3.3500	
	3.3.5200	Updated "interface" parameters
Role	admin	
Example	`switch (config sflow) # agent-ip 20.20.20.20` `switch (config sflow) #`	
Related Commands		
Note	The IP address here is used in the sFlow header.	

clear counters

	clear counters Clears sFlow counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config sflow
History	3.3.3500
Role	admin
Example	`switch (config sflow) # clear counters` `switch (config sflow) #`
Related Commands	
Note	

sflow enable (interface)

	sflow enable **no sflow enable** Enables sFlow on this interface. The no form of the command disables sFlow on the interface.
Syntax Description	N/A
Default	disable no view-port-channel member
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.3.3500 3.3.4500 Added MPO configuration mode
Role	admin

Example	`switch(config  1/1)# sflow enable`
	`...`
	`switch(config interface port-channel 1)# sflow enable`

Related Commands	
Note	

5.22.7.3 Show

show sflow

show sflow

Displays sFlow configuration and counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.3500
	3.6.3004 Updated Example
	3.6.6000 Updated Example
Role	admin

Example	`switch (config)# show sflow`
	`sflow protocol: enabled`
	`sflow: enabled`
	`sampling-rate: 16000`
	`max-sample-size: 128`
	`counter-poll-interval: 20`
	`max-datagram-size: 1400`
	`ip-agent: 0.0.0.0`
	
	`ingress ports:`
	`Interfaces:`
	`Ethernet eth1/2 eth1/1`
	
	`Statistics:`
	`Total Samples: 0`
	`Number of flow samples: 0`
	`Estimated Number of flow discarded: 0`
	`Number of flow statistics samples: 0`
	`Number of datagrams: 0`

Related Commands	
Note	

5.23 802.1x Protocol

The 802.1x standard describes a way to authenticate hosts (or supplicants) and to allow connection only to a list of allowed hosts pre-configured on an authentication server. The authentication is performed by the switch (authenticator) which negotiates the authentication with a RADIUS server (authentication server). This allows to block traffic from non-authenticated sources.

The 802.1x protocol defines the following roles:

- Supplicant – the host. It provides the authentication credentials to the authenticator and awaits approval.
- Authenticator – the device that connects the supplicant to the network, and checks the authentication with the authentication server. The authenticator is also in charge of blocking and isolating of new client till authenticated and allowing communication once the client has passed the authentication. Mellanox switch acts as an authenticator.
- Authentication server – a RADIUS server which can authenticate the user.

The 802.1x is available only on access physical ports. It is not available on LAG and MLAG ports.

A local analyzer port cannot support 802.1x protocol.

802.1x cannot be activated on router port interfaces.

802.1x cannot run on a port configured to switchport trunk or hybrid.

Management interfaces cannot be configured as 802.1x port access entity (PAE) authenticators.

5.23.1 802.1x Operating Modes

The following operating modes are supported in 802.1x:

- Single host – only one supplicant can communicate through the port.

  Once authentication of the supplicant is accepted by the authentication server, the switch allows it access. If the supplicant logs off or the port state is changed, the port becomes unauthenticated. And if a different supplicant tries to access through this port, its bidirectional traffic is discarded (including authentication traffic).

An exception to this is multicast and broadcast traffic which do get transmitted over the interface once authenticated and are exposed to an unauthorized supplicant if it exists.

- Multi-host mode – allows connection of multiple hosts over a single port. Only the first supplicant is authenticated. Subsequent hosts have network access without the need to authenticate.

5.23.2 Configuring 802.1x

➢ ***To configure 802.1x on the switch***

Step 1. Enable 802.1x protocol. Run:

```
switch (config) # protocol dot1x
```

Step 2. Enable the system as authenticator. Run:

```
switch (config) # dot1x system-auth-control
```

Step 3. Configure RADIUS server parameters. Run:

```
switch (config) # dot1x radius-server host 10.10.10.10 key my4uth3nt1c4t10nk3y retrans-
mit 2 timeout 3
```

Step 4. Enter the configuration mode of an Ethernet interface. Run:

```
switch (config) #  1/1
switch (config  1/1) #
```

Step 5. Configure the interface as a port access entity authenticator. Run:

```
switch (config  1/1) # dot1x pae authenticator
```

Step 6. Configure the interface to perform authentication on ingress traffic. Run:

```
switch (config  1/1) # dot1x port-control auto
```

Step 7. Verify 802.1x configuration. Run:

```
switch (config  1/1) # show dot1x interfaces ethernet 1/1

Eth1/1
  PAE Status:                     Enabled
  Configured host mode:           Multi-host
  Configured port-control:        Auto
  Authentication status:          Unauthorized
```

```
  Re-Authentication:               Disabled
  Re-Authentication period (sec):  -
  Tx wait period (sec):            30
  Quiet period (sec):              60
  Max request retry:               2
  Last EAPOL RX source MAC:        00:00:00:00:00:00
switch (config  1/1)#
```

5.23.3 Commands

protocol dot1x

protocol dot1x
no protocol dot1x

Enables 802.1x EAPOL protocol.
The no form of the command disables 802.1x EAPOL protocol.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.2008
Role	admin
Example	`switch (config)# protocol dot1x`
Related Commands	
Note	

dot1x clear-statistics

	dot1x clear-statistics Resets the 802.1x counters on all or a specific port.
Syntax Description	N/A
Default	N/A
Configuration Mode	config config interface ethernet
History	3.4.2008
Role	admin
Example	`switch (config)# dot1x clear-statistics`
Related Commands	
Note	

dot1x pae authenticator

	dot1x pae authenticator **no dot1x pae authenticator** Configures the port as a 802.1x port access entity (PAE) authenticator. The no form of the command disables the port from being a 802.1x PAE authenticator.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet
History	3.4.2008
Role	admin
Example	`switch (config  1/2)# dot1x system-auth-control`
Related Commands	
Note	

dot1x host-mode

dot1x host-mode [multi-host | single-host]
no dot1x host-mode

Configures the authentication mode to either multi-host or single-host.
The no form of the command resets the parameter to its default.

Syntax Description	multi-host	Sets the interface to operate in a port-based mode
	single-host	Sets the interface to operate in a MAC-based mode with support of a single supplicant per interface
Default	single-host	
Configuration Mode	config interface ethernet	
History	3.4.2008	
	3.4.2300	Added “single-host” option
Role	admin	
Example	`switch (config  1/2)# dot1x host-mode single-host`	
Related Commands		
Note		

dot1x port-control

dot1x port-control [auto | force-authorized | force-unauthorized]
no dot1x port-control

Configures 802.1x port access entity (PAE) port-control.
The no form of the command resets the parameter to its default.

Syntax Description	auto	The authenticator uses PAE authentication services to allow or block the port traffic
	force-authorized	Allows traffic on this port regardless of supplicant authorization
	force-unauthorized	Blocks traffic on this port regardless of supplicant authorization
Default	Force-authorized	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2)# dot1x port-control auto`	
Related Commands		
Note		

dot1x radius-server host

dot1x radius-server host <IP address> [enable | auth-port <port> | key <password> | prompt-key | retransmit <retries> | timeout <seconds>]
no dot1x radius-server host <IP address> enable

Configure 802.1x RADIUS server IP address.
The no form of the command disables 802.1x RADIUS server.

Syntax Description	auth-port	Sets 802.1x RADIUS port to use with this server. Range: 1-65535.
	enable	Sets 802.1x RADIUS as administratively enabled
	key	Configures 802.1x global RADIUS shared secret for servers.
	prompt-key	Prompts for key, rather than entering on command line
	retransmit	Configure 802.1x global RADIUS retransmit count for servers. The time configured is in seconds. Range: 0-5.
	timeout	Configures 802.1x global RADIUS timeout value for servers. The time configured is in seconds. Range: 1-60.
Default	auth-port: 1812 key: empty string retransmit: 1 timeout: 3	
Configuration Mode	config	
History	3.4.2008	
Role	admin	
Example	`switch (config)# dot1x radius-server host 10.10.10.10 auth-port 65535 prompt-key enable`	
Related Commands		
Note	• The no form of the various parameters resets them to their default values as indicated in the Default section above • It is possible to configure up to 5 RADIUS servers • It is possible to configure only 1 authentication port per RADIUS server IP	

dot1x reauthenticate

	dot1x reauthenticate **no dot1x reauthenticate** Enables supplicant re-authentication according to the configuration of command "dot1x timeout reauthentication". The no form of the command disables supplicant re-authentication.
Syntax Description	N/A
Default	No re-authentication
Configuration Mode	config interface ethernet
History	3.4.2008
Role	admin
Example	`switch (config  1/2)# dot1x reauthenticate`
Related Commands	
Note	

dot1x system-auth-control

	dot1x system-auth-control **no dot1x system-auth-control** Enables the system as authenticator. The no form of the command disables the system as authenticator.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.2008
Role	admin
Example	`switch (config)# dot1x system-auth-control`
Related Commands	
Note	

dot1x timeout reauthentication

dot1x timeout reauthentication <period>
no dot1x timeout reauthentication

Configures the number of seconds between re-authentication attempts.
The no form of the command resets the parameter to its default.

Syntax Description	period	Time in second. Range: 1-65535 seconds.
Default	3600 seconds	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2)# dot1x timeout reauthentication 3600`	
Related Commands		
Note		

dot1x timeout quiet-period

dot1x timeout quiet-period <period>
no dot1x timeout quiet-period

Configures the number of seconds that the authenticator remains quiet following a failed authentication exchange with the supplicant.
The no form of the command resets the parameter to its default.

Syntax Description	period	Time in second. Range: 1-65535 seconds.
Default	60 seconds	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2)# dot1x timeout quiet-period 60`	
Related Commands		
Note		

dot1x timeout tx-period

dot1x timeout tx-period <period>
no dot1x timeout tx-period

Configures the maximum number of seconds that the authenticator waits for supplicant response of EAP-request/identify frame before retransmitting the request.
The no form of the command resets the parameter to its default.

Syntax Description	period	Time in second. Range: 1-65535 seconds.
Default	30 seconds	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2)# dot1x timeout quiet-period 30`	
Related Commands		
Note		

dot1x max-req

dot1x max-req <retries>
no dot1x max-req

Configures the maximum amount of retries for the authenticator to communicate with the supplicant over EAP.
The no form of the command resets the parameter to its default.

Syntax Description	retries	The number of request retries. Range: 1-10.
Default	2	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2)# dot1x max-req 2`	
Related Commands		
Note		

show dot1x

show dot1x

Displays 802.1x information on all interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Role	admin

Example	`switch (config)# show dot1x`
	
	`System authentication is enabled`
	
	`----------------------------------------------------------------------`
	`Port  Pae  Host-mode  Port-control  Status`
	`----------------------------------------------------------------------`
	`Eth1/1  Enabled  multi-host  auto  unauthorized`
	`Eth1/2  Disabled  multi-host  force-authorized  down`
	`Eth1/3  Disabled  multi-host  force-authorized  down`
	`Eth1/4  Disabled  multi-host  force-authorized  down`
	`Eth1/5  Disabled  multi-host  force-authorized  down`
	`Eth1/6  Disabled  multi-host  force-authorized  down`
	`Eth1/7  Disabled  multi-host  force-authorized  down`
	`Eth1/8  Disabled  multi-host  force-authorized  down`
	`Eth1/9  Disabled  multi-host  force-authorized  down`
	`...`
	`switch (config)#`

Related Commands	
Note	

show dot1x interfaces ethernet

show dot1x interfaces ethernet <slot>/<port>

Displays 802.1x interface information.

Syntax Description	<slot>/<port> Ethernet interface
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Role	admin

Example	`switch (config)# show dot1x interfaces ethernet 1/2`
	
	`Eth1/2`
	`  PAE Status:  Enabled`
	`  Configured host mode:  Multi-host`
	`  Configured port-control:  Auto`
	`  Authentication status:  Unauthorized`
	`  Re-Authentication:  Enabled`
	`  Re-Authentication period (sec):  3600`
	`  Tx wait period (sec):  30`
	`  Quiet period (sec):  60`
	`  Max request retry:  2`
	`  Last EAPOL RX source MAC:  00:00:00:00:00:00`
	`switch (config  1/2)#`

Related Commands	
Note	

show dot1x interfaces ethernet statistics

show dot1x interfaces ethernet <slot>/<port> statistics

Displays 802.1x interface information.

Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
Role	admin	

Example	`switch (config)# show dot1x interfaces ethernet 1/2 statistics`
	
	`Eth1/2`
	`  EAPOL frames received: 3`
	`  EAPOL frames transmitted: 2`
	`  EAPOL Start frames received: 1`
	`  EAPOL Logoff frames received: 0`
	`  EAP Response-ID frames received: 2`
	`  EAP Response frames received: 0`
	`  EAP Request-ID frames transmitted: 2`
	`  EAP Request frames transmitted: 0`
	`  Invalid EAPOL frames received: 0`
	`  EAP length error frames received: 0`
	`  Last EAPOL frame version: 1`
	`  Last EAPOL frame source: 00:1A:A0:02:E9:8E`
	`switch (config)#`	

Related Commands		
Note		

show dot1x radius

	show dot1x radius Displays 802.1x RADIUS settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Role	admin

Example	`switch (config)# show dot1x radius`
	`802.1x RADIUS defaults:`
	`  Key:  ********`
	`  Timeout:  3`
	`  Retransmit:  1`
	`No 802.1x RADIUS servers configured.`
	`switch (config)#`

Related Commands	
Note	

5.24 Voice VLAN

Voice VLAN allows configuring a port to provide QoS to voice and data traffic in a scenario where a terminal is connected to an IP phone which is in turn connected to the port on the switch. The IP phone bridges the data traffic from the terminal into the switch port. Any voice traffic from the IP phone is also sent to the same port with no differentiation. Therefore it is in the administrator's interest to provide different QoS to the voice traffic and the data traffic by placing the voice traffic on a different VLAN from the data traffic.

This can be achieved by configuring a voice VLAN on the desired switch port using LLDP-MED TLVs. Media Endpoint Discovery (MED) TLVs allow the switch to apply certain policies by informing the remote media device to configure itself using different TLV.

In this use-case scenario we employ the use of the network policy TLV, which is defined as per TIA-TR41. The network policy TLV can be used to inform a specific VLAN to use for an application stream.

Onyx allow the user to configure the VLAN for voice traffic. In Figure 32, the user configures a voice VLAN of 25 and the switch port has a PVID of 50. Therefore all the voice traffic is switched onto VLAN 25 and the untagged packets from the terminal are switched into VLAN 50.

Figure 32: Tagging Voice Packets with a Different VLAN ID

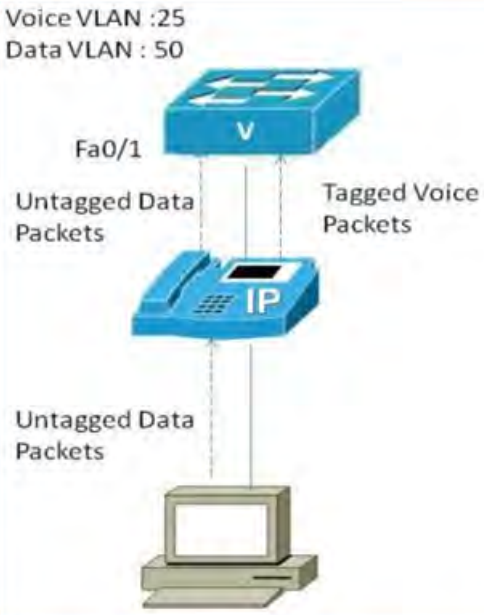

5.24.1 Configuring Voice VLAN

- ***To configure LLDP-MED TLV, run:***

```
switch (config) #  1/4
switch (config  1/4) # lldp med-tlv-select media-capabilities
switch (config  1/4) # lldp med-tlv-select network-policy
switch (config  1/4) # lldp med-tlv-select all
```

- ***To verify LLDP-MED TLV configuration, run:***

```
switch (config) # show lldp interfaces
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC:
Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------
Eth1/1    Enabled   Enabled   PD, SD
Eth1/2    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/3    Disabled  Disabled  PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-NWP
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R,
                              MED-CAP, MED-NWP
Eth1/5    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/6    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
...

switch (config) # show lldp interfaces ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC:
Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R,
                              MED-CAP, MED-NWP

switch (config) # show lldp interfaces ethernet 1/4 med-cap
Media Capabilities:
    LLDP-MED Capab   : Yes
    Network Policy   : Yes
    Location Id      : No
    Ext Power MDI-PSE: No
    Ext Power MDI-PD : No

Network Policy:
    Application Type : 1 (Voice)
    VLAN Id          : 11
    L2 Priority      : 0
    DSCP Value       : 0
```

➢ ***To configure voice VLAN:***

Step 1. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # exit
switch (config) #
```

Step 2. Set the interface mode to be hybrid. Run:

```
switch (config) #  1/4 switchport mode hybrid
switch (config) #  1/4 switchport hybrid allowed-vlan 200
```

Step 3. Assign the VLAN to the interface. Run:

```
switch (config) #  1/4 switchport voice vlan 200
```

Step 4. (Optional) Change the PVID of the port so that untagged packets go to a different VLAN than the default. Run:

```
switch (config)# vlan 300
switch (config vlan 300)# exit
switch (config)#  1/4 switchport access vlan 300
```

Step 5. Verify the configuration. Run:

```
switch (config)# show interfaces switchport
Interface       Mode         Access vlan        Allowed vlans
-------------------------------------------------------------------------------
Eth1/1          access       1
Eth1/2          access       1
Eth1/3          access       1
Eth1/4          hybrid       300                200
Eth1/5          access       1
...
switch (config)# show lldp interfaces ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC:
Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R,
                              MED-CAP, MED-NWP
switch (config)# show lldp interfaces ethernet 1/4 med-cap
Media Capabilities:
    LLDP-MED Capab   : Yes
    Network Policy   : Yes
    Location Id      : No
    Ext Power MDI-PSE: No
    Ext Power MDI-PD : No
```

```
Network Policy:
    Application Type : 1 (Voice)
    VLAN Id          : 200
    L2 Priority      : 0
    DSCP Value       : 0
```

➢ ***To remove voice VLAN and LLDP-MED TLV:***

Step 1. Remove the voice VLAN from the interface. Run:

```
switch (config)# no  1/4 switchport voice vlan
```

Step 2. Disable the MED TLV from the interface. Run:

```
switch (config)#  1/4 lldp med-tlv-select none
```

5.24.2 Limitations

1. LLDP MED cannot be enabled on a router port interface and vice versa (i.e. a port that has LLDP MED enabled cannot be configured as a router port interface).
2. LLDP MED cannot be enabled on a LAG and vice versa (i.e. a port that has LLDP MED enabled cannot be configured as a LAG).
3. If switchport is in trunk, dot1q-tunnel, or dcbx-access, configuring either the TLV or Voice VLAN gives a warning message.

5.25 Store-and-Forward

5.25.1 General

Store-and-Forward is used to describe a functionality where a switch receives a complete packet, stores it, and only then forwards it.

since the switch make forwarding decisions based on the destination address which is at the header of the packet, the switch can make the forwarding decision before receiving the complete packet, this process is called cut-through, the switch forwards part of the packet before receiving the complete packet.

Cut-through allows lower latency and saves buffer space, but if an error occurred in the packet while utilizing cut-through, the packet will be forwarded with an error, alternatively, utilizing store-and-forward allows the switch to drop erroneous packets.

The standard implementation of forwarding mode is for the entire switch; either all ports on a switch are in store-and-forward mode or all ports on a switch are in cut-through mode.

Mellanox implements forwarding mode per egress port, which is a more flexible method and vital in cases where a switch is connected to both a storage device and a compute server among other setups.

5.25.2 Commands

switchmode store-and-forward

[no | disable] switchmode store-and-forward

Enables global store-and-forward configuration on the switch.
The no form of the command removes store-and-forward configuration from the switch and reverts it back to the switch's global configuration.
The disable form of the command configures the forwarding mode to cut-through.

Syntax Description	N/A
Default	N/A
Configuration Mode	config config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.3640 3.6.6000 — Added "config interface mlag port channel" configuration mode
Role	admin
Example	`switch (config)# switchmode store-and-forward`
Related Commands	
Note	

6 IP Routing

6.1 General

6.1.1 IP Interfaces

Mellanox Onyx™ supports the following 3 types of IP interfaces:

- VLAN interface
- Loopback interface
- Router port interface

Onyx supports up to 999 IP interfaces.

Each IP interface can be configured with multiple IP addresses. The first address assigned to the interface automatically becomes its primary address (only one primary address is supported per interface), and the rest are secondary addresses.

Secondary addresses are advertised via OSPF. No “HELLO” messages are sent on them and no adjacencies are established on them either.

Primary addresses cannot be modified once assigned. To assign a different primary address, all addresses of the interface must be removed and then reconfigured.

Up to 16 IPv4 (as well as IPv6) addresses are supported on each IP interface.

6.1.1.1 VLAN Interfaces

VLAN interface is a logical IPv4 interface created per subnet over a specific 802.1Q VLAN ID. If two hosts from two different subnets need to communicate (via the IP layer), the network administrator needs to configure two interface VLANs, one for each of the subnets.

Each interface VLAN has the following attributes:

- Admin state
- Operational state
- MAC address
- IP address and mask
- MTU
- Description
- Set of counters

6.1.1.2 Loopback Interfaces

Loopback interface is a logical software entity where traffic transmitted to this interface is immediately received on the sending end.

6.1.1.3 Router Port Interfaces

Router port interface is a regular switch port configured to operate as an L3 interface. Router port interfaces are assigned an IP address and all L3 commands become applicable to them.

Once configured, router port interfaces no longer partake in the bridging activities of the switch and VLANs configured on them are separate from the pool allocated for the switch ports.

6.1.1.4 Configuring a VLAN Interface

➢ ***To configure a VLAN interface:***

Step 1. Create a VLAN. Run:

```
switch (config)# vlan 10
switch (config vlan 10)# exit
```

Step 2. Assign a physical interface to this VLAN. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport mode access
switch (config interface ethernet 1/1)# exit
```

Step 3. There must be at least one interface in the operational state “UP”.

```
switch (config)# show  1/1 status
Port                 Operational state          Speed                Negotiation
----                 -----------------          -----                -----------
Eth1/1               Up                        40 Gbps              No-Negotiation
```

Step 4. Create a VLAN interface that matches the VLAN. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10)#
```

Step 5. Configure an IP address and a network mask to the interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.10 /24
```

Step 6. Verify VLAN interface configuration. Run:

```
switch (config interface vlan 10) # show interfaces vlan 10

Vlan 10:
  Admin state      : Enabled
  Operational state: Down
  Autostate        : Enabled
  Mac Address      : 24:8A:07:F3:04:C8
  DHCP client      : Disabled

  IPv4 address:
    10.10.10.10/24 [primary]

  Broadcast address:
    10.10.10.255 [primary]

  Arp responder: Disabled
  MTU          : 1500 bytes
  Arp timeout  : 1500 seconds
  Icmp redirect: Enabled
  Description  : my-ip-interface
  VRF          : default
  Counters     : Disabled
```

6.1.1.5 Configuring a Loopback Interface

➢ ***To configure a loopback interface:***

Step 1. Create a loopback interface. Run:

```
switch (config)# interface loopback 2
switch (config interface loopback 2)#
```

Step 2. Configure an IP address on the loopback interface. Run:

```
switch (config interface loopback 2)# ip address 20.20.20.20 /32
```

Step 3. Verify loopback interface configuration. Run:

```
switch (config interface loopback 2)# show interfaces loopback 2

Loopback 2:
  IPv4 address:
    20.20.20.20/32 [primary]

  Broadcast address:
    20.20.20.20 [primary]

  MTU        : 1500 bytes
  Description: my-loopback
  VRF        : default
```

6.1.1.6 Configuring a Router Port Interface

Step 1. Enter an Ethernet interface's configuration context. Run:

```
switch (config)# interface ethernet 1/10
switch (config interface ethernet 1/10)#
```

Step 2. Configure the Ethernet interface to become an router port interface. Run:

```
switch (config interface ethernet 1/10)# no switchport force
```

Step 3. Configure an IP address on the router port interface. Run:

```
switch (config interface ethernet 1/10)# ip address 100.100.100.100 /24
```

Step 4. Verify router port interface configuration. Run:

```
switch (config interface ethernet 1/10)# show interfaces ethernet 1/10

Eth1/10:
  Admin state                     : Enabled
  Operational state               : Down
  Last change in operational status: Never
  Boot delay time                 : 0 sec
  Description                     : N/A
  Mac address                     : 24:8A:07:F3:04:C8
  MTU                             : 1500 bytes (Maximum packet size 1522 bytes)
  Fec                             : auto
  Flow-control                    : receive off send off
  Supported speeds                : 1G 10G 25G
  Advertised speeds               : 1G 10G 25G
  Actual speed                    : Unknown
  Auto-negotiation                : Enabled
  Width reduction mode            : Unknown
  DHCP client                     : Disabled
  Autoconfig                      : Disabled

  IPv4 address:
    100.100.100.100/24 [primary]

  Broadcast address:
    100.100.100.255 [primary]

  Arp responder: Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                 : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                  : 0 bits/sec, 0 bytes/sec, 0 packets/sec
```

```
Rx:
  0                    packets
  0                    unicast packets
  0                    multicast packets
  0                    broadcast packets
  0                    bytes
  0                    discard packets
  0                    error packets
  0                    fcs errors
  0                    undersize packets
  0                    oversize packets
  0                    pause packets
  0                    unknown control opcode
  0                    symbol errors

Tx:
  0                    packets
  0                    unicast packets
  0                    multicast packets
  0                    broadcast packets
  0                    bytes
  0                    discard packets
  0                    error packets
  0                    hoq discard packets
```

6.1.2 Equal Cost Multi-Path Routing (ECMP)

Equal-cost multi-path routing (ECMP) is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple paths.

In Figure 33, routers R1 and R2 can both access each of their router peer networks. Router R1 routing table for 10.0.40/24 will contain the following routes:

- 10.0.10.2
- 10.0.20.2
- 10.0.30.2

Figure 33: ECMP

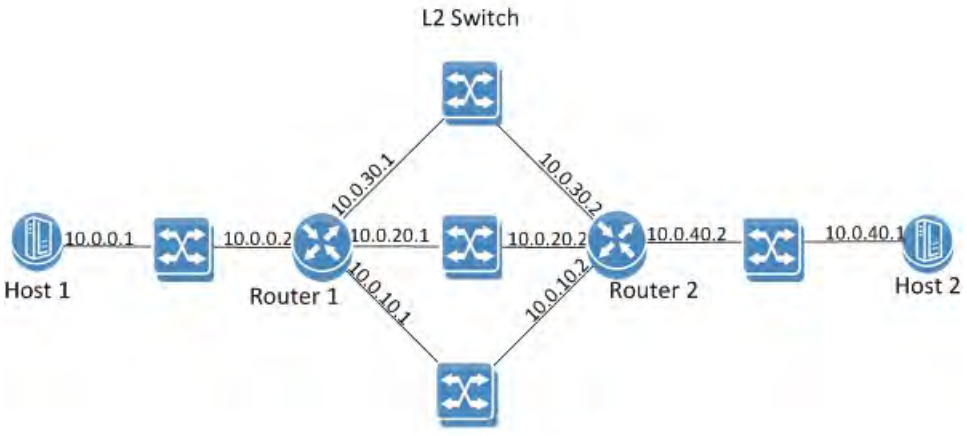

The load balancing function of the ECMP is configured globally on the system.

Hash algorithm can be symmetric or asymmetric. In symmetric hash functions bidirectional flows between routes will follow the same path, while in asymmetric hash functions, bidirectional traffic can follow different paths in both directions.

The following load balancing types are supported:

- Source IP & Port – source IP (SIP) and source UDP/TCP port: If the packet is not UDP/TCP, only SIP is used for the hash calculation. This is an asymmetric hash function.
- Destination IP & Port – destination IP (DIP) and destination UDP/TCP port: If the packet is not UDP/TCP, only DIP is used for the hash calculation. This is an asymmetric hash function.
- Source and Destination IP & Port – destination and source IP, as well as destination and source UDP/TCP port: If the packet is not UDP/TCP, only SIP/DIP are used for the hash calculation. This is a symmetric hash function.
- Traffic Class – Load balance based on the traffic class assigned to the packet. This is an asymmetric hash function.
- All (default) – all above fields are part of the hash calculations. This is a symmetric hash function.

6.1.2.1 Hash Functions

It is advised that LAG and ECMP hash function configuration over more than one hop is different. If the same hash function is used over two hops, all the traffic sorted from one hop to following one will arrive already having the same characteristics, which will render the next hash function useless. For example, configure load-balancing on the first hop based on source IP while on the next hop based on destination IP.

Figure 34: Multiple Hash Functions

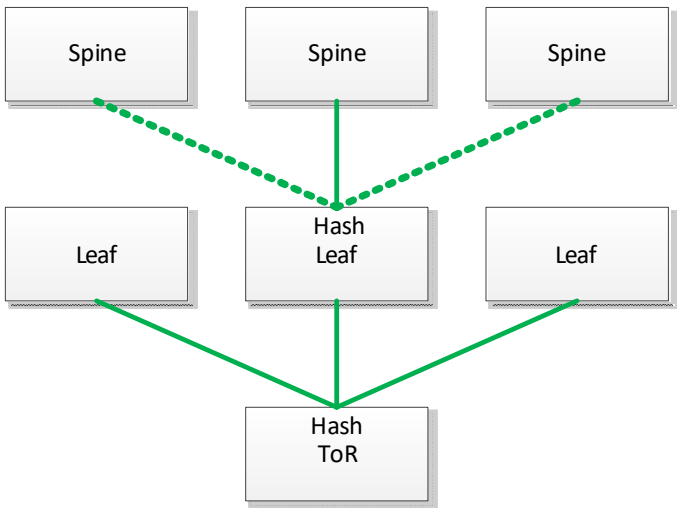

6.1.2.2 ECMP Consistent Hashing

In an IP network multiple flows share the same path defined by their destination prefix. ECMP allows those flows to travel with the same prefix and be distributed over multiple next hops that usually belong to different physical links, in order to reach better bandwidth utilization. When using the standard ECMP some links in the network become unreachable, thus the next hop list and hash function distribution change, and flows are moved to other links. Packet reordering in the network or failure in a user session might occur, while others which use anycast IP addresses utilize ECMP distribution for load balancing. Therefore, changing the next hop may cause flows to arrive to the wrong destination.

When network is reconfigured, and route next hop set is changed, flows that are not affected by the change should continue to be sent to the same next hops and keep the same outgoing link.

Using consistent hash containers enables you to use size arrays with next hop buckets to make sure unaffected flows are sent to the same next hops when some next hops are removed from the container. When a new next hop is added to the consistent hash container, some buckets are replaced with a new next hop, so part of the existing flows are moved to a new next hop.

When a route is installed, it points to a hash container. Each flow in the route is mapped to a respective bucket, and is eventually forwarded to the next hop in the bucket.

In the following example we see a single route with 3 flows and 4 next hops, so the container has 12 bucket.

Figure 35: Consistent Hashing #1

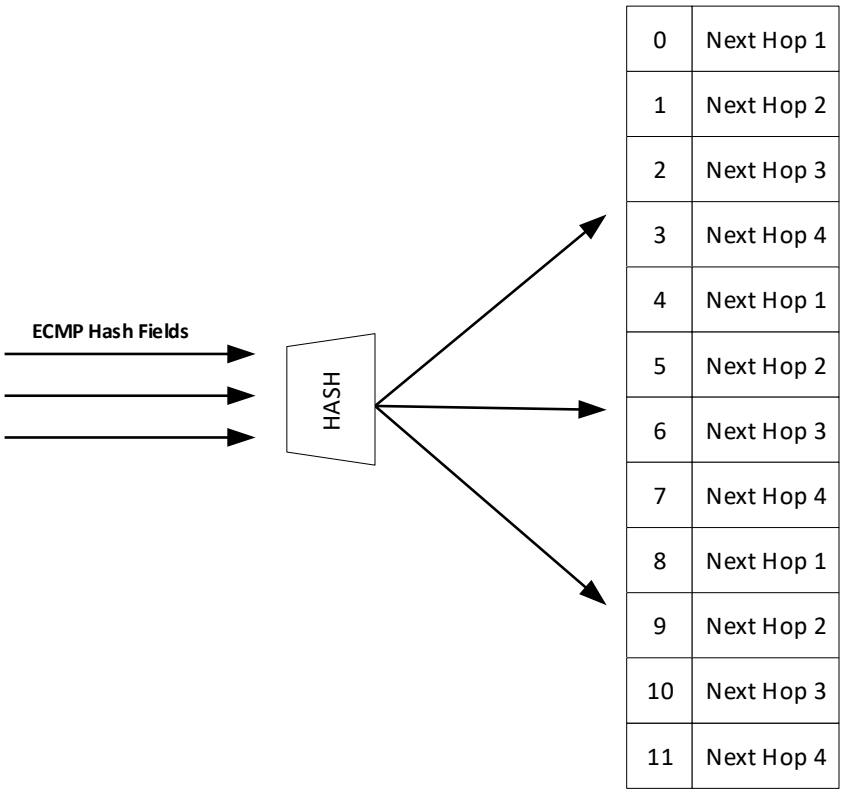

6.1.2.2.1 Remove Next Hops

Unlike the default IP load-sharing hashing, when consistent hashing is used, and a next hop needs to be removed, the number of hash buckets does not change. All appearances of the deleted next hop are removed from the container and replaced by the remaining next hops.

Figure 36: Consistent Hashing #2

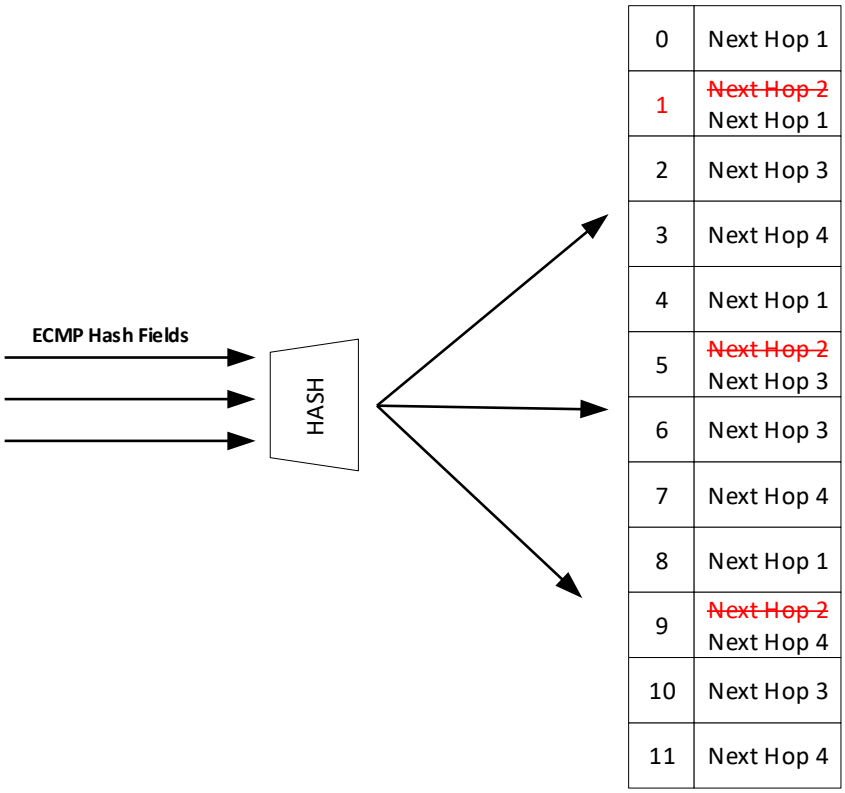

6.1.2.2.2 Add Next Hops

When adding a new next hop, some of existing next hops should be removed from the hash, and the new next hop should be located in one of the newly available places.The new next hops are not applied to HW immediately, but only after a convergence time period.

Figure 37: Consistent Hashing #3

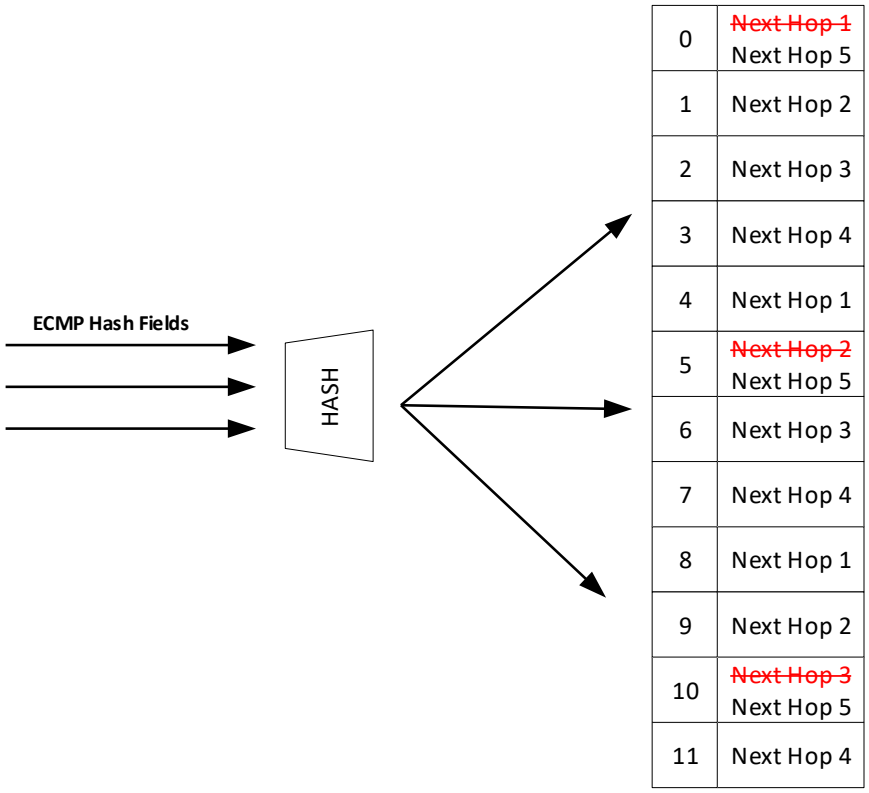

6.1.2.2.3 Supported Number of Containers

Please refer to the following information:

Container bucket size	Default Number of Containers	Maximum Number of Containers
512	40	96
1024	20	48

When the consistent hashing containers count exceeds the maximum number of containers, the operational state of consistent hashing function will become 'unstable' and the containers with the same next hop sets will be merged to release more resources. Once more resources are available to deploy the containers, the operation state will become 'stable'.

In the unstable case which may result from lack of consistent hashing resources, the new route will be installed as a non-consistent route, and a random next hop from its next hop set will be

chosen as the actual next hop and installed in hardware. The route will only be partially programed in hardware.

6.1.2.2.4 Configuring Consistent Hashing

To configure consistent hashing, run “ip load-sharing type consistent.

6.1.2.3 Virtual Routing and Forwarding

Virtual Routing and Forwarding (VRF) allows multiple routing table instances to coexist within the same router simultaneously. Since the routing instances are independent, IP addresses on each routing table may overlap without conflicting with each other.

VRF can be used for the following purposes:

- Ensure customer privacy and security
- Separate between management and user data
- Support customers with the same address space
- Support VPN

Multiple routing instances defined in the router can have different purposes and can be configured in different manners:

- Different IP interfaces can be attached to different VRFs (only one IP interface can be in a single VRF)
- Routing in VRF can be enabled or disabled
- Each VRF component can run its own routing protocol independently from other instances
- Differently configured IPv4 and IPv6 services

The first VRF in the system is created automatically and it is called “default” VRF. It cannot be deleted or configured.

Onyx supports up to 64 VRFs, 8 instances of BGP, and 8 instances of OSPF.

6.1.3 ARP Neighbor Discovery Responder

ARP functionality in IP/Ethernet networks is needed to provide mapping from IP addresses to L2 MAC addresses. This request may be sent in multiple cases:

- A station wants to initiate an IP session with another station on the same IP subnet and needs to obtain its L2 address
- A station wants to update other stations that its MAC address has changed
- A station wants to check that the MAC address of its peer did not change
- The peer responds with unicast ARP response.

The following are two scenarios when ARP responder functionality is needed:

- Network wants to avoid broadcast in the network or on some parts of the network, so broadcast ARP packets are not distributed in that part of the network

- There is no L2 connectivity between some parts of the network, and even IP addressing scheme does not reflect it

ARP responder answers a broadcast ARP requests that arrive to the switch.

ARP responder is configured on an IP interface (with or without IP address) of any type (e.g. VLAN interface, router port, or LAG).

Only IP interfaces in UP admin state respond to ARP.

This functionality is provided for all ARP entries that are configured or provided on the interface: Static, dynamic, or per protocol.

There is no need to enable IP routing in the system to enable ARP responder functionality.

If a user has multiple VRFs the interface can be created in any VRF. If IP routing is disabled the interface is created in default VRF.

ARP responder can be enabled together with IP routing and given an interface which can be used in routing.

When IP routing on the interface is enabled, all entries that have been used by the responder become ARP entries for the router and vice versa.

A user must avoid using ARP responder in broadcast networks—the system itself does not block it.

6.1.3.1 Configuring ARP Responder

➢ ***In order to initialize ARP responder:***

Step 1. Create IP interface. Run:

```
switch (config) # interface vlan 10
switch (config interface vlan 10) #
```

Step 2. Initialize ARP responder on the interface. Run:

```
switch (config interface vlan 10) # ip arp responder
```

Step 3. Create static ARP entries on VLAN. Run:

```
switch (config interface vlan 10) # ip arp 172.130.11.1 00:11:22:33:44:55
```

Step 4. Create ACL to drop broadcast, and assign it to all relevant L2 interface (VLAN's members). Run:

```
switch (config interface vlan 10) # mac access-list new
switch (config interface vlan 10) # mac access-list new seq-number 10 deny
any FF:FF:FF:FF:FF:FF mask FF:FF:FF:FF:FF:FF
switch (config interface vlan 10) #  1/3-1/5 mac port access-group new
```

6.1.4 Commands

6.1.4.1 General

ip l3

ip l3 [force]
no ip l3 [force]

Enables IP routing capabilities.
The no form of the command disables IP routing and removes its configuration.

Syntax Description	N/A
Default	L3
Configuration Mode	config
History	3.4.1802
Role	admin
Example	`switch (config) # ip l3 force`
Related Commands	N/A
Note	

vrf definition

vrf definition <vrf-name>

Creates the VRF.

Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
	3.6.6000	Updated the notes section
Role	admin	
Example	`switch (config) # vrf definition my-vrf` `switch (config vrf definition my-vrf) #`	
Related Commands	N/A	
Notes	63 VRFs are supported aside from the default VRF	

routing-context vrf

routing-context vrf <vrf-name>

Enters the active-context of the specified session.

Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
Role	admin	
Example	`switch (config) # routing-context vrf my-vrf` `switch (config) #`	
Related Commands	N/A	
Notes	• If a routing-context is configured, the user does not have to explicitly specify the VRF name parameter in this or any other VRF command • If no routing-context is configured and the user does not specify the VRF name, default VRF is used	

ip routing

ip routing [vrf <vrf-name>]

Enables L3 forwarding between high speed interfaces.

Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config	
History	3.4.1802	
	3.4.2008	Added VRF parameter
Role	admin	
Example	`switch (config) # ip routing vrf my-vrf` `switch (config) #`	
Related Commands	N/A	
Notes	• RD must be configured to enable IP routing on the VRF • If no routing-context is specified, the “routing-context” VRF is automatically configured.	

description

description <description>
no description force

Creates the VRF.

Syntax Description	description	Text string
	force	Forces deletion (no confirmation needed if configuration exists inside the VRF)
Default	N/A	
Configuration Mode	config vrf definition	
History	3.4.2008	
Role	admin	
Example	`switch (config vrf definition my-vrf) # description vrf-description` `switch (config vrf definition my-vrf) #`	
Related Commands	N/A	
Notes		

rd

rd [<ip addr>:<0-65,535> | <AS Number>:<0-4,294,967,295> | <AS Number>:<ip addr>]

Adds a route distinguisher (RD) to the VRF configuration mode.

Syntax Description	ip-addr	IPv4 address
	AS Number	Asynchronous machine number
Default	N/A	
Configuration Mode	config vrf definition	
History	3.4.2008	
Role	admin	
Example	`switch (config vrf definition my-vrf) # rd 10.10.10.10:2` `switch (config vrf definition my-vrf) #`	
Related Commands	N/A	
Notes	• RDs internally identify routes belonging to a VRF to distinguish overlapping or duplicate IP address ranges. This allows the creation of distinct routes to the same IP address for different VPNs. The RD is a 64-bit number made up of an AS number or IPv4 address followed by a user-selected ID number. Once an RD has been assigned to a VRF it cannot be changed. To change the RD, remove the VRF then create it again. VRF is not active until an RD is defined. • An RD must be defined to enable IP routing on the VRF	

vrf forwarding

vrf forwarding <vrf-name>

Maps an interface to VRF.

Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config interface ethernet set as router port interface config interface vlan config interface loopback	
History	3.4.2008	
Role	admin	
Example	`switch (config  1/2) # vrf forwarding my-vrf` `switch (config  1/2) #`	
Related Commands	N/A	
Notes		

clear ip routing counters

	clear ip routing counters Clears counters, related to NULL interface and dropped packets by router.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6102
Role	admin
Example	`switch (config) # clear ip routing counters`
Related Commands	N/A
Notes	

show ip routing

show ip routing [vrf <vrf-name> | all]

Displays IP routing information per VRF.

Syntax Description	vrf	Displays information for specific VRF
	all	Displays information on all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.4.2008	Added VRF parameter
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show ip routing`
	
	`VRF Name default:`
	`  IP routing: enabled`
	
	`switch (config) # show ip routing vrf all`
	
	`VRF Name default:`
	`  IP routing: enabled`
	
	`VRF Name new:`
	`  IP routing: disabled`	

Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

show ip routing counters

show ip routing [vrf <vrf-name> | all] counters

Display counters, related to NULL interface and dropped packets by router.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6102
Role	admin

Example	`switch (config) # show ip routing counters`
	`  1  packets discarded by router`
	`  64  bytes discarded by router`
	`  2  packets to null interface`
	`  128  bytes to null interface`

Related Commands	N/A
Notes	

s

show routing-context vrf

	show routing-context vrf Displays VRF active context.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Role	admin

Example	`switch (config) # show routing-context vrf`
	`VRF active context: my-vrf`
	`switch (config) #`

Related Commands	N/A
Notes	

show vrf

show vrf [<vrf-name> | all]

Displays VRF information.

Syntax Description	all	Displays information for all VRF instances
	vrf-name	Name of VRF instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.6000	Updated Example
Role	admin	

Example	`switch (config) # show vrf my-vrf`
	`VRF Info:`
	`  Name: default`
	`  RD: NA`
	`  Description: NA`
	`  IP routing state: Disabled`
	`  IPv6 routing state: Disabled`
	`  IP multicast routing state: Disabled`
	`  Protocols:`
	`  Interfaces:`	

Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

6.1.4.2 IP Interfaces

switchport

switchport [force]
no switchport [force]

Configures the Ethernet interface as a regular switchport.
The no form of the command configures the Ethernet interface as router port interface.

Syntax Description	force	Forces configuration even if the interface's admin state is enabled.
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.3.5200	
	3.6.4006	Added storm-control support
Role	admin	

Example	`switch (config  1/10)# no switchport force`
	`error message is case storm-control is configured on port:`
	`% interface * has storm control configuration. Please remove it first`	

Related Commands		
Note	When storm-control is configured on port, an err-msg will appear. Force command deletes all storm-control configuration from port.	

encapsulation dot1q vlan

encapsulation dot1q vlan <vlan-id> [force]
no encapsulation dot1q vlan [force]

Enables L2 802.1Q encapsulation of traffic on a specified router port interface in a VLAN.
The no form of the command disables L2 802.1Q encapsulation of traffic on a specified router port interface in a VLAN.

Syntax Description	vlan-id	Enables L2 802.1Q encapsulation of traffic on a router port interface in a VLAN
	force	Forces admin state down
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.3.5200	
Role	admin	
Example	`switch (config  1/10)# encapsulation dot1q vlan 10`	
Related Commands		
Note		

6.1.4.3 Interface VLAN

interface vlan

interface vlan <vlan-id>
no interface vlan <vlan-id>

Creates a VLAN interface and enters the interface VLAN configuration mode.
The no form of the command deletes the VLAN interface.

Syntax Description	vlan-id	A numeric range of 1-4094
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
Role	admin	
Example	`switch (config) # interface vlan 10` `switch (config interface vlan 10) #`	
Related Commands	ip routing vlan <vlan-id> switchport mode switchport access show interfaces vlan	
Note	• Make sure the VLAN was created, using the command “vlan <vlan-id>” in the global configuration mode • The VLAN must be assigned to one of the L2 interfaces. To do so, run the command “swichport ...” • At least one interface belong to that VLAN must be in UP state	

interface vlan <id> no-autostate

[no] interface vlan <id> no-autostate

Disables the VLAN interface autostate such that its operational state remains up as long as its admin state is up, even if no port in the relevant VLAN egress-list is operationally up.

Syntax Description	vlan-id	A numeric range of 1-4094 or a range of VLANs.
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Role	admin	
Example	`switch (config) # interface vlan 10 no-autostate` `switch (config) # interface vlan 10-13 no-autostate`	
Related Commands	show ip interface vlan <id>	
Note		

ip address

ip address <ip-address> <mask>
no ip address [<ip-address> [<mask>]]

Enters user-defined IPv4 address for the interface.
The no form of the command removes the specified IPv4 address. If no address is specified, then all IPv4 addresses of this interface are removed.

Syntax Description	ip-address	IPv4 address
	mask	There are two possible ways to the mask: • /length (i.e. /24) • Network address (i.e. 255.255.255.0) The mask length may be configured without a space (i.e. <ipv4-address>/<length>)
Default	0.0.0.0/0	
Configuration Mode	config interface vlan	
History	3.2.0230	
Role	admin	
Example	`switch (config interface vlan 10) # ip address 10.10.10.10 /24`	
Related Commands	interface vlan show interfaces vlan	
Note	An interface may have up to 16 IPv4 address assignments	

counters

	counters **no counters** Enables counters on the IP interface. The no form of the command disables counters gathering on the IP interface.
Syntax Description	N/A
Default	counters are disabled.
Configuration Mode	config interface vlan
History	3.2.0230
Role	admin
Example	`switch (config interface vlan 10) # counters` `switch (config interface vlan 10) #`
Related Commands	counters interface vlan show interfaces vlan
Note	• Enabling counters for the router interface adds delay to the traffic stream • There are maximum of 16 counter sets

description

description <string>
no description

Enters a description for the interface.
The no form of the command sets the description to default.

Syntax Description	string	User defined string
Default	“”	
Configuration Mode	config interface vlan	
History	3.2.0230	
Role	admin	
Example	`switch (config interface vlan 10) # description my-ip-interface` `switch (config interface vlan 10) #`	
Related Commands	interface vlan show interfaces vlan	
Note		

mtu

mtu <size> [force]
no mtu

Sets the MTU for the interface.
The no form of the command sets the MTU to default.

Syntax Description	size	1500-9216.
	force	Forces command implementation.
Default	1522	
Configuration Mode	config interface vlan	
History	3.2.0230	
Role	admin	
Example	`switch (config interface vlan 10)# mtu 9216` `switch (config interface vlan 10 #`	
Related Commands	interface vlan show interfaces vlan	
Note		

shutdown

	shutdown **no shutdown** Disables the interface. The no form of the command enables the interface.
Syntax Description	N/A
Default	The interface is enabled.
Configuration Mode	config interface vlan
History	3.1.0000
Role	admin
Example	`switch (config interface vlan 20) # shutdown` `switch (config interface vlan 20) #`
Related Commands	interface vlan
Note	

clear counters

	clear counters Clears the interface counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface vlan
History	3.2.0230
Role	admin
Example	`switch (config interface vlan 10) # clear counters` `switch (config interface vlan 10) #`
Related Commands	interface vlan counters
Note	

ip icmp redirect

	ip icmp redirect **no ip icmp redirect** Enables ICMP redirect. The no form of the command disables ICMP redirect.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface vlan
History	3.4.0010
Role	admin
Example	`switch (config interface vlan 10) # no ip icmp redirect`
Related Commands	interface vlan counters
Note	• ICMP redirect transmits messages to hosts alerting them about the existence of more efficient routes to a specific destination

show interfaces

show interfaces [brief]

Displays interface configuration.

Syntax Description	brief	Displays brief output
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.3000	
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show interfaces`
	
	`Interface lo status:`
	`  Comment  :`
	`  Admin up  : yes`
	`  Link up  : yes`
	`  DHCP running  : no`
	`...`
	`Interface mgmt0 status:`
	`  Comment  :`
	`  Admin up  : yes`
	`  Link up  : yes`
	`  DHCP running  : yes`
	`...`
	`Interface mgmt1 status:`
	`  Comment  :`
	`  Admin up  : yes`
	`  Link up  : yes`
	`  DHCP running  : yes (but no valid lease)`
	`...`
	`Eth1/1:`
	`  Admin state  : Enabled`
	`  Operational state  : Up`
	`  Last change in operational status: 0:22:11 ago (5 oper change)`
	`  Boot delay time  : 0 sec`
	`...`	

Related Commands	N/A	
Notes		

show interfaces vlan

show interfaces vlan [<id>]

Displays interface configuration.

Syntax Description	id	Specifies the VLAN ID for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.3000	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show interfaces vlan 100

Vlan 100:
  Admin state      : Enabled
  Operational state: Down
  Autostate        : Enabled
  Mac Address      : 24:8A:07:83:30:C8
  DHCP client      : Disabled

  IPv4 address:
    192.168.70.254/24 [primary]
    192.168.80.254/24

  Broadcast address:
    192.168.70.255 [primary]
    192.168.80.255

  IPv6 address:
    4000::1/64 [primary]
    5000::1/64

  MTU          : 1500 bytes
  Arp timeout  : 1500 seconds
  Icmp redirect: Enabled
  Description  : N/A
  VRF          : default
  Counters     : Disabled
```

Related Commands	N/A
Notes	

show ip interface

show ip interface [vrf <vrf-name>]

Displays IP interfaces information.

Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip interface

Interface mgmt0 status:
  Comment         :
  Admin up        : yes
  Link up         : yes
  DHCP running    : yes
...
Interface mgmt1 status:
  Comment         :
  Admin up        : yes
  Link up         : yes
  DHCP running    : yes (but no valid lease)
...
Vlan 100:
  Admin state      : Enabled
  Operational state: Down
  Autostate        : Enabled
  Mac Address      : 24:8A:07:83:30:C8
...
Eth1/1:
  Admin state                     : Enabled
  Operational state               : Up
  Last change in operational status: 0:14:39 ago (5 oper change)
  Boot delay time                 : 0 sec
...
Po1:
  Admin state         : Enabled
  Operational state   : Down
  Description         : N\A
  Mac address         : 24:8A:07:83:30:C8
...
Loopback 1:
  IPv4 address:
    192.168.1.1/32 [primary]
    192.168.2.1/32
...
```

Related Commands	N/A
Notes	

show ip interface brief

show ip interface [vrf <vrf-name] brief

Displays IP interfaces brief information.

Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip interface brief

-------------------------------------------------------------------------------------------------
Interface     Address/Mask          Primary      Admin-state      Oper-state      MTU       VRF
-------------------------------------------------------------------------------------------------
mgmt0         10.12.67.33/25                     Enabled          Up              1500      default
mgmt1         Unassigned                         Enabled          Up              1500      default
Vlan 100      192.168.70.254/24     primary      Enabled          Down            1500      default
Vlan 100      192.168.80.254/24
Eth1/1        192.168.50.254/24     primary      Enabled          Up              1500      default
Eth1/1        192.168.60.254/24
Po1           192.168.100.254/24    primary      Enabled          Down            1500      default
Po1           192.168.110.254/24
Loopback 1    192.168.1.1/32        primary      Enabled          Up              1500      default
Loopback 1    192.168.2.1/32
```

Related Commands	N/A
Notes	

show interface configured

show ip interface [<type> <id>] configured

Displays interface configuration.

Syntax Description	<type> <id>	Specifies the interface for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show interfaces mgmt0 configured`
	
	`Interface mgmt0 configuration:`
	`  Comment  :`
	`  Enabled  : yes`
	`  DHCP  : yes`
	`  DHCP Hostname  : yes`
	`  Zeroconf  : no`
	`  IP address  :`
	`  Netmask  :`
	`  IPv6 enabled  : yes`
	`  Autoconf enabled: no`
	`  Autoconf route  : yes`
	`  Autoconf privacy: no`
	`  DHCPv6 enabled  : yes`
	`  IPv6 addresses  : 0`
	`  Speed  : auto`
	`  Duplex  : auto`
	`  MTU  : 1500`	

Related Commands	N/A	
Notes		

show ip

show ip interface [vrf <vrf-name>] ethernet <slot>/<port>

Displays information on the specified Ethernet interface in the routing-context VRF.

Syntax Description	<slot>/<port>	Port number
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
Role	admin	
Example		

```
switch (config) # show ip  1/1

Eth1/1:
  Admin state                     : Enabled
  Operational state               : Up
  Last change in operational status: 0:11:14 ago (5 oper change)
  Boot delay time                 : 0 sec
  Description                     : N/A
  Mac address                     : 24:8A:07:83:30:C8
  MTU                             : 1500 bytes (Maximum packet size 1522 bytes)
  Fec                             : auto
  Flow-control                    : receive off send off
  Supported speeds                : 1G 10G 25G
  Advertised speeds               : 1G 10G 25G
  Actual speed                    : 25G (auto)
  Auto-negotiation                : Enabled
  Width reduction mode            : Unknown
  DHCP client                     : Disabled
  Autoconfig                      : Disabled

  IPv4 address:
    192.168.50.254/24 [primary]
    192.168.60.254/24

  Broadcast address:
    192.168.50.255 [primary]
    192.168.60.255

  IPv6 address:
    2000::1/64 [primary]
    3000::1/64
    fe80::268a:7ff:fe83:30c8/64

  Arp responder  : Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through
```

```
Telemetry sampling: Disabled   TCs: N\A
Telemetry threshold: Disabled  TCs: N\A
Telemetry threshold level: N\A

Last clearing of "show interface" counters: Never
60 seconds ingress rate                  : 56 bits/sec, 7 bytes/sec, 1 packets/sec
60 seconds egress rate                   : 8 bits/sec, 1 bytes/sec, 0 packets/sec

Rx:
   698                  packets
   0                    unicast packets
   0                    multicast packets
   698                  broadcast packets
   44672                bytes
   0                    discard packets
   0                    error packets
   0                    fcs errors
   0                    undersize packets
   0                    oversize packets
   0                    pause packets
   0                    unknown control opcode
   0                    symbol errors
Tx:
   1923                 packets
   0                    unicast packets
   1859                 multicast packets
   64                   broadcast packets
   142718               bytes
   0                    discard packets
   0                    error packets
   0                    hoq discard packets
```

Related Commands N/A

Notes

show ip interface mgmt0

show ip interface [vrf <vrf-name>] mgmt0

Displays management interface information.

Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip interface mgmt0

Interface mgmt0 status:
  Comment         :
  Admin up        : yes
  Link up         : yes
  DHCP running    : yes
  IP address      : 10.12.67.33
  Netmask         : 255.255.255.128
  IPv6 enabled    : yes
  Autoconf enabled: no
  Autoconf route  : yes
  Autoconf privacy: no
  DHCPv6 running  : yes (but no valid lease)
  IPv6 addresses  : 1

  IPv6 address:
    fe80::268a:7ff:fe53:3d8e/64

  Speed           : 1000Mb/s (auto)
  Duplex          : full (auto)
  Interface type  : ethernet
  Interface source: bridge
  MTU             : 1500
  HW address      : 24:8A:07:53:3D:8E

  Rx:
    1843422  bytes
      25627  packets
          0  mcast packets
          0  discards
          0  errors
          0  overruns
          0  frame
```

```
Tx:
  236174  bytes
    1897  packets
       0  discards
       0  errors
       0  overruns
       0  carrier
       0  collisions
       0  queue len
```

Related Commands	N/A
Notes	

show ip interface port-channel

show ip interface [vrf <vrf-name>] port-channel <id>

Displays information on the specified LAG in the routing-context VRF.

Syntax Description	id	LAG ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated Example
	3.7.100x	Updated Example
Role	admin	
Example		

```
switch (config) # show ip interface port-channel 1

Po1:
  Admin state         : Enabled
  Operational state   : Down
  Description         : N/A
  Mac address         : 24:8A:07:83:30:C8
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : 25G (auto)
  Auto-negotiation    : N/A
  Width reduction mode: Not supported
  DHCP client         : Disabled
  Autoconfig          : Disabled

  IPv4 address:
    192.168.100.254/24 [primary]
    192.168.110.254/24

  Broadcast address:
    192.168.100.255 [primary]
    192.168.110.255

  IPv6 address:
    6000::1/64 [primary]
    7000::1/64

  Arp responder  : Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A
```

```
Last clearing of "show interface" counters: Never
60 seconds ingress rate                  : 0 bits/sec, 0 bytes/sec, 0 packets/sec
60 seconds egress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec

Rx:
  0                    packets
  0                    unicast packets
  0                    multicast packets
  0                    broadcast packets
  0                    bytes
  0                    discard packets
  0                    error packets
  0                    fcs errors
  0                    undersize packets
  0                    oversize packets
  0                    pause packets
  0                    unknown control opcode
  0                    symbol errors

Tx:
  0                    packets
  0                    unicast packets
  0                    multicast packets
  0                    broadcast packets
  0                    bytes
  0                    discard packets
  0                    error packets
  0                    hoq discard packets
```

Related Commands	N/A
Notes	

show ip interface vrf

show ip interface vrf {<vrf-name> | all | ethernet <slot>/<port> | loopback <id> | port-channel <id> | vlan <vid>} [brief]

Displays IP interface information per VRF.

Syntax Description	vrf	Displays IP interface information per VRF
	all	Displays information on all VRF
	ethernet	Displays Ethernet interface information per VRF
	loopback	Displays loopback interface information per VRF
	port-channel	Displays LAG information per VRF
	vlan	Displays VLAN interface information per VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.5000	Updated Example
	3.6.6000	Updated Example
	3.6.8008	Updated Example
	3.7.100x	Updated Example
Role	admin	

Example

```
switch (config) # show ip interface vrf default port-channel 1

Po1:
  Admin state         : Enabled
  Operational state   : Down
  Description         : N/A
  Mac address         : 24:8A:07:83:30:C8
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : 25G (auto)
  Auto-negotiation    : N/A
  Width reduction mode: Not supported
  DHCP client         : Disabled
  Autoconfig          : Disabled
...
```

Related Commands	N/A

Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.

show ipv6 interface

show ipv6 interface

Displays IPv6 interface information.

Syntax Description	vrf VRF name
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Role	admin
Example	

```
switch (config) # show ipv6 interface

Eth1/1:
  VRF        : default
  Admin state: enabled
  IPv6       : enabled

  IPv6 address:
    2000::1/64 [primary]
    3000::1/64

  Local Link Address:
    fe80::268a:7ff:fe83:30c8/64

  Joined group address:
    ff02::1:ff00:1

  ND retransmit interval (usec): 1000
  ND DAD                       : enabled
  Number of DAD attempts       : 1
  ND reachable time            : 0

Po1:
  VRF        : default
  Admin state: enabled
  IPv6       : enabled

  IPv6 address:
    6000::1/64 [primary]
    7000::1/64

ND retransmit interval (usec): 1000
  ND DAD                       : enabled
  Number of DAD attempts       : 1
  ND reachable time            : 0

vlan100:
  VRF        : default
  Admin state: enabled
  IPv6       : enabled
```

```
 IPv6 address:
    4000::1/64 [primary]
    5000::1/64

  ICMPv6 redirect              : enabled
  ND retransmit interval (usec): 1000
  ND DAD                       : enabled
  Number of DAD attempts       : 1
  ND reachable time            : 0

loopback1:
  VRF        : default
  Admin state: enabled
  IPv6       : enabled

  IPv6 address:
    2001::1/128 [primary]
    2002::1/128

  Local Link Address:
    fe80::4c01:40ff:feb3:b753/64

  Joined group address:
    ff02::1:ff00:1
```

Related Commands	N/A
Notes	

show ipv6 interface brief

show ipv6 interface [vrf <vrf-name>] brief

Displays IPv6 interface information.

Syntax Description	vrf VRF name
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Role	admin

Example

```
switch (config) # show ipv6 interface brief
-------------------------------------------------------------------------------------------------------------------
Interface     Address/Mask                    Primary      Address-state     Admin-state     Oper-state     MTU      VRF
-------------------------------------------------------------------------------------------------------------------
mgmt0         fe80::268a:7ff:fe53:3d8e/64                  valid             Enabled         Up             1500     default
mgmt1         fe80::268a:7ff:fe53:3d8f/64                  valid             Enabled         Up             1500     default
Eth1/1        2000::1/64                      primary      valid             Enabled         Up             1500     default
Eth1/1        3000::1/64                                   valid
Eth1/1        fe80::268a:7ff:fe83:30c8/64                  valid
Po1           6000::1/64                      primary      valid             Enabled         Down           1500     default
Po1           7000::1/64                                   valid
vlan100       4000::1/64                      primary      valid             Enabled         Down           1500     default
vlan100       5000::1/64                                   valid
loopback1     2001::1/128                     primary      valid             Enabled         Up             1500     default
loopback1     2002::1/128                                  valid
loopback1     fe80::4c01:40ff:feb3:b753/64                 valid
```

Related Commands	N/A
Notes	

show ipv6

show ipv6 interface [vrf <vrf-name>] ethernet <slot>/<port>

Display IPv6 information of the specified Ethernet interface.

Syntax Description	<slot>/<port>	Port number
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ipv6  1/1`
	
	`Eth1/1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    2000::1/64 [primary]`
	`    3000::1/64`
	
	`  Local Link Address:`
	`    fe80::268a:7ff:fe83:30c8/64`
	
	`  Joined group address:`
	`    ff02::1:ff00:1`
	
	`  ND retransmit interval (usec): 1000`
	`  ND DAD                       : enabled`
	`  Number of DAD attempts       : 1`
	`  ND reachable time            : 0`	

Related Commands	N/A	
Notes		

show ipv6 interface loopback

show ipv6 interface [vrf <vrf-name>] loopback <id>

Display IPv6 information of the specified loopback interface.

Syntax Description	id	Loopback port ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ipv6 interface loopback 1`
	
	`loopback1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    2001::1/128 [primary]`
	`    2002::1/128`
	
	`  Local Link Address:`
	`    fe80::4c01:40ff:feb3:b753/64`
	
	`  Joined group address:`
	`    ff02::1:ff00:1`	

Related Commands	N/A	
Notes		

show ipv6 interface port-channel

show ipv6 interface [vrf <vrf-name>] port-channel <id>

Display IPv6 information of the specified LAG interface.

Syntax Description	id	LAG ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ipv6 interface port-channel 1`
	
	`Po1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    6000::1/64 [primary]`
	`    7000::1/64`
	
	`  ND retransmit interval (usec): 1000`
	`  ND DAD                       : enabled`
	`  Number of DAD attempts       : 1`
	`  ND reachable time            : 0`	

Related Commands	N/A	
Notes		

show ipv6 interface vlan

show ipv6 interface [vrf <vrf-name>] vlan <vid>

Display IPv6 information of the specified VLAN interface.

Syntax Description	vid	VLAN ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ipv6 interface vlan 100`
	
	`vlan100:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    4000::1/64 [primary]`
	`    5000::1/64`
	
	`  ICMPv6 redirect              : disabled`
	`  ND retransmit interval (usec): 1000`
	`  ND DAD                       : enabled`
	`  Number of DAD attempts       : 1`
	`  ND reachable time            : 0`	

Related Commands	N/A	
Notes		

show ipv6 interface vrf

show ipv6 interface vrf <vrf-name>

Display IPv6 information of the specified VRF.

Syntax Description	name	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Role	admin	

Example	`switch (config) # show ipv6 interface vrf default`
	
	`Eth1/1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	`...`
	`Po1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	`...`
	`vlan100:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	`...`
	`loopback1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	`...`	

Related Commands	N/A	
Notes		

show ipv6 interface vrf brief

show ipv6 interface vrf <name> brief

Display IPv6 information of the specified VRF in brief form.

Syntax Description	name VRF name
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Role	admin

Example

```
switch (config) # show ipv6 interface vrf default brief
-------------------------------------------------------------------------------------------------------------------------
Interface     Address/Mask                    Primary      Address-state     Admin-state     Oper-state     MTU      VRF
-------------------------------------------------------------------------------------------------------------------------
mgmt0         fe80::268a:7ff:fe53:3d8e/64                  valid             Enabled         Up             1500     default
mgmt1         fe80::268a:7ff:fe53:3d8f/64                  valid             Enabled         Up             1500     default
Eth1/1        2000::1/64                      primary      valid             Enabled         Up             1500     default
Eth1/1        3000::1/64                                   valid
Eth1/1        fe80::268a:7ff:fe83:30c8/64                  valid
Po1           6000::1/64                      primary      valid             Enabled         Down           1500     default
Po1           7000::1/64                                   valid
vlan100       4000::1/64                      primary      valid             Enabled         Down           1500     default
vlan100       5000::1/64                                   valid
loopback1     2001::1/128                     primary      valid             Enabled         Up             1500     default
loopback1     2002::1/128                                  valid
loopback1     fe80::4c01:40ff:feb3:b753/64                 valid
```

Related Commands	N/A
Notes	

6.1.4.4 Loopback Interface

interface loopback

interface loopback <id>
no interface loopback <id>

Creates a loopback interface and enters the interface configuration mode.
The no form of the command deletes the interface.

Syntax Description	id	A numeric range of 0-31
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Role	admin	
Example	`switch (config) # interface loopback 10` `switch (config interface loopback 10) #`	
Related Commands		
Note	• Up to 32 loopback interfaces can be configured • Within the loopback configuration mode, you can configure description and ip-address • MTU cannot be configured on the loopback interface	

ip address

ip address <ip-address> <mask>
no ip address [<ip-address> [<mask>]]

Enters user-defined IPv4 address for the interface.
The no form of the command removes the specified IPv4 address. If no address is specified, then all IPv4 addresses of this interface are removed.

Syntax Description	ip-address	IPv4 address
	mask	There are two possible ways to the mask: • /length – only /32 is possible • Network address (i.e. 255.255.255.0) The mask length may be configured without a space (i.e. <ipv4-address>/<length>).
Default	0.0.0.0/0	
Configuration Mode	config interface loopback	
History	3.3.5006	
Role	admin	
Example	`switch (config interface loopback 10) # ip address 10.10.10.10 /32`	
Related Commands	interface loopback	
Note	An interface may have up to 16 IPv4 address assignments.	

description

description <string>
no description

Enters a description for the interface.
The no form of the command sets the description to default.

Syntax Description	string	User defined string.
Default	“”	
Configuration Mode	config interface loopback	
History	3.3.5006	
Role	admin	
Example	`switch (config interface loopback 10) # description my-ip-interface`	
Related Commands	interface loopback	
Note		

show interfaces loopback

show interface loopback <id>

Shows the attribute of the interface loopback.

Syntax Description	id	A numeric range of 1-32
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config) # show interfaces loopback 1`
	
	`Loopback 1:`
	`  IPv4 address:`
	`  192.168.1.1/32 [primary]`
	`  192.168.2.1/32`
	
	`  Broadcast address:`
	`  192.168.1.1 [primary]`
	`  192.168.2.1`
	
	`  IPv6 address:`
	`  2001::1/128 [primary]`
	`  2002::1/128`
	`  fe80::4c01:40ff:feb3:b753/64`
	
	`  MTU  : 1500 bytes`
	`  Description: N/A`
	`  VRF  : default`	

Related Commands		
Note		

6.1.4.5 Routing and ECMP

ip route

ip route [vrf <vrf-name>] <ip-prefix> <netmask> <next -hop-ip-address> [<distance>]
no ip route [vrf <vrf-name>] <ip-prefix> <netmask> [<next -hop-ip-address>]

Configures a static route inside VRF.
The no form of the command removes the static route configured.

Syntax Description	vrf-name	VRF session name
	ip-prefix	IP address
	netmask	There are two possible ways to input the mask: • /<length> (e.g. /24) • Network address (e.g. 255.255.255.0)
	next-hop-ip-address	IP address of the next hop
	distance	Administrative distance assigned to route. Options include: • No parameter – route is assigned a default administrative distance of 1 • 1-255 – the administrative distance assigned to route
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
	3.6.6000	Removed ethernet, port-channel, and vlan parameters
Role	admin	
Example	`switch (config) # ip route vrf my-vrf 80.80.80.0 /24 20.20.20.2`	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

ip load-sharing

ip load-sharing <type> [ecmp-group-size <size> [max-ecmp-groups <max>]]
no ip load-sharing

This command sets the ECMP load sharing mode.
The no form of the command sets the load-sharing to default.

Syntax Description	type	• source-ip-port – source ip and TCP/UDP port • destination-ip-port – destination ip and TCP/UDP port • source-destination-ip-port – source & destination ip and TCP/UDP port • flow-label – flow label • all – all options • consistent- consistent hashing mode
	ecmp-group-size	Configures ECMP consistent hashing group size
	max-ecmp-groups	Configures max groups of ECMP consistent hashing
Default	all	
Configuration Mode	config	
History	3.7.11xx	Updated syntax
	3.2.0230	
	3.5.1000	Added flow-label parameter
Role	admin	
Example	`switch (config) # ip load-sharing all` `switch (config) # ip load-sharing consistent [ecmp-group-size<size>]`	
Related Commands	ip route	
Note		

show ip route

show ip route [vrf <vrf-name] [[<ip-address> | <ip-address>/<length>] [longer-prefixes]] [connected | bgp | static]

Displays routing table.

Syntax Description	ip-address	Performs longest prefix match (LPM) and displays best route
	<ip-address>/<length>	Displays next hop for the specified network. If the network does not exist in routing table, it is not shown. Note: It is the user's responsibility to calculate the mask and enter it correctly. For example: • Valid – show ip route 10.10.10.0/24 • Invalid – show ip route 10.10.10.10/24
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination. (Only available if both IP and mask are specified.)
	connected	Displays entries for routes to networks directly connected to the switch
	bgp	Display BGP routes
	static	Displays entries added through CLI commands
History	3.7.11xx	Updated Example
	3.6.5000	Updated Example
	3.6.6000	Updated Example
	3.6.8008	Updated Example
Role	admin	
Example		

```
switch (config) # show ip route

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)
  c: consistent hashing
  p: partial programming in H/W

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  default         0.0.0.0                 10.12.67.126    mgmt0          DHCP     1/1
  10.12.67.0      255.255.255.128         0.0.0.0         mgmt0          direct   0/0
  192.168.2.0     255.255.255.0     c     0.0.0.0         vlan1          direct   0/0
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

show ip route vrf

show ip route vrf {<vrf-name> | all}

Displays routing table of VRF instance.

Syntax Description	all	Displays routing tables for all VRF instances
	vrf-name	Name of VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.4070	Added support for BFD and updated notes
	3.6.5000	Updated Example
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip route vrf default

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  default         0.0.0.0                 10.12.67.126    mgmt0          DHCP     1/1
  10.12.67.0      255.255.255.128         0.0.0.0         mgmt0          direct   0/0

switch (config) # show ip route vrf my-vrf static

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name my-vrf:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  80.80.80.0      255.255.255.0           20.20.20.2      vlan20         static   1/1
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed
	• When using a network prefix, the user must calculate the host mask and enter correctly. For example, “show ip route 10.10.10.0/24” is valid, but “ip route 10.10.10.10/24” is invalid.

show ip route -a

show ip route [vrf {<vrf-name> | all}] -a

Displays routing table of VRF instance.

Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
	-a	Displays static routes currently inactive due to the interface being down
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Role	admin	

Example	`switch (config) # show ip route vrf my-vrf -a`
	`VRF Name: my-vrf`
	`---`
	`Destination Mask Gateway Interface Source Distance/Metric`
	`90.90.90.0 255.255.255.0 1.1.1.2 NA static 1/0`	

Related Commands	ip route	

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed	

show ip route failed

show ip route [vrf {<vrf-name> | all}] failed

Displays failed routes of VRF instance.

Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip route failed

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

Warning: Number of HW failed routes is 2
These routes are marked with 'f' flag

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  20.20.20.0      255.255.255.0   f      0.0.0.0         vlan20         direct   0/0
  80.80.80.0      255.255.255.0   f      20.20.20.2      vlan20         static   1/1
```

Related Commands

ip route

Notes

- If no default route exists, then the message “Route not found” is printed
- Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
  - If route is removed from routing decision it will be marked as “Active”
  - Protected next hops are marked with “B”
  - BFD protected failed/non active neighbors are marked with “BF”
- If no routing-context is specified, the “routing-context” VRF is automatically displayed

show ip route static

show ip route [vrf {<vrf-name> | all}] static

Displays static routes of VRF instance.

Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.5000	Updated Example
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip route static

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  80.80.80.0      255.255.255.0          20.20.20.2      vlan20         static   1/1
```

Related Commands

ip route

Notes

- If no default route exists, then the message “Route not found” is printed
- Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
  - If route is removed from routing decision it will be marked as “Active”
  - Protected next hops are marked with “B”
  - BFD protected failed/non active neighbors are marked with “BF”
- If no routing-context is specified, the “routing-context” VRF is automatically displayed

show ip route static multicast-override

show ip route [vrf {all | <vrf-name>}] static multicast-override

Displays Reverse Path Forwarding (RPF) information for a specific IPv4 multicast source configured via the command "ip mroute".

Syntax Description	vrf-name	Name of VRF
	all	Displays information for all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.8008	Updated Example
Role	admin	
Example	`switch (config) # show ip route vrf default static multicast-override`	

```
VRF "default":
  ------------------------------------------------------------------------
  Destination       Mask              Gateway           Route preference
  ------------------------------------------------------------------------
  50.50.50.0        255.255.255.0     20.20.20.45       1
  100.100.8.0       255.255.255.0     20.20.20.9        1
  100.100.100.0     255.255.255.0     20.20.20.22       7
  100.100.100.100   255.255.255.255   20.20.20.9        1
```

Related Commands	
Notes	

show ip route summary

show ip route [vrf {<vrf-name> | all}] summary

Displays route summary of VRF instance.

Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.5000	Updated Example
Role	admin	

Example	`switch (config) # show ip route vrf my-vrf summary`
	`VRF Name:  default`
	
	`------------------------`
	`Route Source  Routes`
	`------------------------`
	`direct  3`
	`static  0`
	`ospf  0`
	`bgp  0`
	`DHCP  1`
	`Total  4`	

Related Commands	ip route	

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed	

show ip route interface

show ip route [vrf {<vrf-name> | all}] interface {ethernet <slot>/<port> | port-channel <lag> | vlan <vlan>}

Displays routing table for specific interfaces.

Syntax Description	ethernet	Displays routing table for Ethernet interfaces
	port-channel	Displays routing table for LAG interfaces
	vlan	Displays routing table for VLAN interfaces
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	Added VRF parameter
	3.6.5000	Updated Example
Role	admin	
Example		
Related Commands	ip route	
Notes		

```
switch (config) # show ip route interface vlan 10
VRF Name:         default
Total number of entries: 1

------------------------------------------------------------------------------
  Address             Type           Hardware Address          Interface
------------------------------------------------------------------------------
  15.0.0.2            Static ETH     DE:DE:BE:EF:DE:AD         vlan 10
```

show ip load-sharing

show ip load-sharing

Displays ECMP hash attribute.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.7.11xx Updated example
	3.2.0230
Role	admin

Example	`(config) # show ip load-sharing`
	`Load sharing: all`
	`Type: static`
	
	`(config) # show ip load-sharing`
	`Load sharing: destination-ip-port`
	`Type: consistent`
	`Operational state: stable`
	`Container size: 512`
	`Max number of containers: 40`
	`Used containers: 5`

Related Commands	ip load-sharing
Note	The command's output is different for static & consistent hashing

6.1.4.6 Network to Media Resolution (ARP)

ip arp

ip arp [vrf <vrf-name>] <ip-address> <mac-address>
no ip arp <ip-address>

Configures IP ARP properties of VRF
The no form of the command deletes the static ARP configuration.

Syntax Description	vrf-name	VRF session name
	IP address	IPv4 address
	mac-address	MAC address (format XX:XX:XX:XX:XX:XX)
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
Role	admin	
Example	`switch (config) # ip arp vrf my-vrf 20.20.20.2 aa:bb:cc:dd:ee:ff`	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

ip arp responder

	ip arp responder Initiates ARP responder functionality.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ethernet config interface port-channel config interface vlan
History	3.6.8008
Role	admin
Example	`switch (config interface vlan 10) # ip arp responder`
Related Commands	ip arp show ip arp
Note	

ip arp timeout

ip arp timeout <timeout-value>
no ip arp timeout

Sets the dynamic ARP cache timeout.
The no form of the command sets the timeout to default.

Syntax Description	timeout-value	Time (in seconds) that an entry remains in the ARP cache. Range: 240-28800.
Default	1500 seconds	
Configuration Mode	config interface ethernet config interface vlan config interface port-channel	
History	3.2.0230	
	3.5.1000	Updated Note section
Role	admin	
Example	`switch (config) # ip arp timeout 2000`	
Related Commands	ip arp show ip arp	
Note	• This configuration may take up to 5 minutes to take effect • The time interval after which each ARP entry becomes stale may actually vary from 50-150% of the configured value	

clear ip arp

clear ip arp [vrf <vrf-name>] [interface <type> | <IP-address>]

Clears the dynamic ARP cache for the specific VRF session.

Syntax Description	vrf-name	VRF session name
	interface	Clears dynamic ARP entries for a interface
	ip-address	Clears dynamic ARP entries for a specific IP address
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
History	3.4.2008	Added VRF parameter
Role	admin	
Example	`switch (config) # clear ip arp vrf my-vrf`	
Related Commands	ip arp show ip arp	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

show ip arp

show ip arp [vrf [<vrf-name> | all]] [interface <type> | count]

Displays all ARP information for VRF instance.

Syntax Description	all	Displays all ARP information for all VRF
	interface	Displays all ARP information for specific interface
	count	Displays number of ARPs for specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3000	
	3.4.2008	Added VRF parameter
	3.6.5000	Updated example output
Role	admin	

Example	`switch (config) # show ip arp vrf my-vrf interface vlan 20`
	`VRF Name: default`
	`Total number of entries: 1`
	
	`---------------------------------------------------------------------------`
	`Address Type Hardware Address Interface`
	`---------------------------------------------------------------------------`
	`15.0.0.2 Static ETH DE:DE:BE:EF:DE:AD vlan 10`	

Related Commands	ip arp	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

6.1.4.7 IP Diagnostic Tools

ping

ping [vrf <vrf-name>] [-LRUbdfnqrvVaA] [-c count] [-i interval] [-w deadline] [-p pattern] [-s packetsize] [-t ttl] [-I interface or address] [-M mtu discovery hint] [-S sndbuf] [-T timestamp option] [-Q tos] [hop1 ...] destination

Sends ICMP echo requests to a specified host.

Syntax Description	Linux Ping options	
	vrf	Specifies VRF instance name
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Role	admin	

Example	`switch (config) # ping 172.30.2.2`
	`PING 172.30.2.2 (172.30.2.2) 56(84) bytes of data.`
	`64 bytes from 172.30.2.2: icmp_seq=1 ttl=64 time=0.703 ms`
	`64 bytes from 172.30.2.2: icmp_seq=2 ttl=64 time=0.187 ms`
	`64 bytes from 172.30.2.2: icmp_seq=3 ttl=64 time=0.166 ms`
	`64 bytes from 172.30.2.2: icmp_seq=4 ttl=64 time=0.161 ms`
	`64 bytes from 172.30.2.2: icmp_seq=5 ttl=64 time=0.153 ms`
	`64 bytes from 172.30.2.2: icmp_seq=6 ttl=64 time=0.144 ms`
	`^C`
	`--- 172.30.2.2 ping statistics ---`
	`6 packets transmitted, 6 received, 0% packet loss, time 5004ms`
	`rtt min/avg/max/mdev = 0.144/0.252/0.703/0.202 ms`	

Related Commands	traceroute	
Note	When using -I option use the interface name + interface number, for example “ping -I vlan10”	

traceroute

traceroute [vrf <vrf-name>] [-46dFITUnrAV] [-f first_ttl] [-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] host [packetlen]

Traces the route packets take to a destination.

Syntax Description		
	vrf	Specifies VRF instance name
	-4	Uses IPv4.
	-6	Uses IPv6
	-d	Enables socket level debugging.
	-F	Sets DF (“do not fragment” bit) on.
	-I	Uses ICMP ECHO for tracerouting.
	-T	Uses TCP SYN for tracerouting.
	-U	Uses UDP datagram (default) for tracerouting.
	-n	Does not resolve IP addresses to their domain names.
	-r	Bypasses the normal routing and send directly to a host on an attached network.
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses.
	-V	Prints version info and exit.
	-f	Starts from the first_ttl hop (instead from 1).
	-g	Routes packets throw the specified gateway (maximum 8 for IPv4 and 127 for IPv6).
	-i	Specifies a network interface to operate with.
	-m	Sets the max number of hops (max TTL to be reached). Default is 30.
	-N	Sets the number of probes to be tried simultaneously (default is 16).
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets.

-l	Uses specified flow_label for IPv6 packets.
-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
-q	Sets the number of probes per each hop. Default is 3.
-s	Uses source src_addr for outgoing packets.
-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).

Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Role	admin	

Example	`switch (config) # traceroute 192.168.10.70`
	`traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte packets`
	`1 172.30.0.1 (172.30.0.1) 3.632 ms 2.849 ms 3.544 ms`
	`2 10.222.128.46 (10.222.128.46) 3.176 ms 3.289 ms 3.656 ms`
	`3 10.158.128.30 (10.158.128.30) 15.331 ms 15.819 ms 16.388 ms`
	`4 10.158.128.65 (10.158.128.65) 20.468 ms 7.893 ms 12.27 ms`
	`5 10.7.34.115 (10.7.34.115) 16.405 ms 11.985 ms 12.264 ms`
	`6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms`	

Related Commands		
Note	• The following flags are not supported: -6, -l, -A • When using -i option use the interface name + interface number, for example “traceroute -i vlan10”	

tcpdump

tcpdump [vrf <vrf-name>] [-aAdeflLnNOpqRStuUvxX] [-c count] [-C file_size]
[-E algo:secret] [-F file] [-i interface] [-M secret]
[-r file] [-s snaplen] [-T type] [-w file]
[-W filecount] [-y datalinktype] [-Z user]
[expression]

Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.

Syntax Description	vrf	Specifies VRF instance name
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Role	admin	

Example	`switch (config) # tcpdump`
	`......`
	`09:37:38.678812 IP 192.168.10.7.ssh > 192.168.10.1.54155: P`
	`1494624:1494800(176) ack 625 win 90`
	`<nop,nop,timestamp 5842763 858672398>`
	`09:37:38.678860 IP 192.168.10.7.ssh > 192.168.10.1.54155: P`
	`1494800:1495104(304) ack 625 win 90`
	`<nop,nop,timestamp 5842763 858672398>`
	`...`
	`9141 packets captured`
	`9142 packets received by filter`
	`0 packets dropped by kernel`
	`switch (config) #`	

Related Commands	N/A	
Note	• When using -i option use the interface name + interface number, for example “tcpdump -i vlan10” • For all flag options of this command refer to the linux ‘man page’ of tcp dump.	

6.1.4.8 QoS

qos map dscp-to-pcp preserve-pcp

	qos map dscp-to-pcp preserve-pcp **no qos map dscp-to-pcp preserve-pcp** Configures the router to copy PCP bits when transferring data from one subnet to another. The no form of the command disables this ability.
Syntax Description	N/A
Default	Disabled.
Configuration Mode	config
History	3.3.4000
Role	admin
Example	`switch (config) # qos map dscp-to-pcp preserve-pcp`
Related Commands	
Note	

6.2 IPv6

IP version 6 (IPv6) is a routing protocol which succeeds IPv4. With the expansion of the Internet and data bases IPv6 addresses consist of 128 bits whose purpose is to allow networks to include a significantly higher number of nodes by increasing the pool of available unique IP addresses. IPv6 packets alleviate overhead and allow for future customizability.

Textual representations of IPv6 addresses consist of 128 bits made up from eight 16-bit hexadecimal numbers separated by colons. IPv6 addresses may be abbreviated as follows:

- You may omit leading zeros in each 16-bit sequence
- You may replace an entire sequence with a double colon if it equals zero

For example, these addresses represent the same IPv6 address:

- af23:0000:0000:0000:1284:037d:35ce:2401
- af23:0:0:0:1284:37d:35ce:2401
- af23::1284:37d:35ce:2401

IPv6 addresses typically denote a 64-bit network prefix and a 64-bit host address.

6.2.1 Neighbor Discovery Protocol

Neighbor Discovery (ND) decides relationships between neighbors and replaces ARP, ICMP, and ICMP redirect in IPv4.

Five kinds of ICMPv6 packets are defined by ND:

- Neighbor advertisement
- Router advertisement
- Neighbor solicitation
- Router solicitation
- Redirect

ND checks whether a neighboring node’s address has changed, whether the neighbor is still reachable, and also resolves the address of the neighbor which a packet is being forwarded to. ND is also useful for network nodes for discovering other nodes and performing basic link-layer configuration.

6.2.2 Configuring IPv6

Figure 38: IPv6 Network

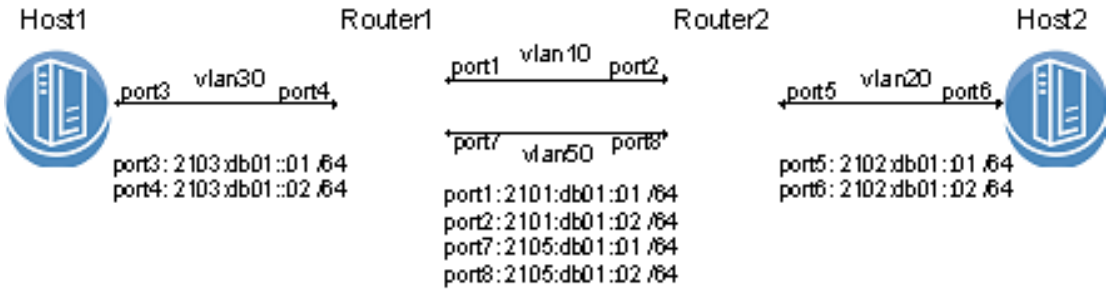

➢ ***To configure Router1:***

Step 1. Enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

Step 3. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 30
switch (config interface vlan 30) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

Step 4. Enable IPv6 on the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 enable
switch (config)# interface vlan 30 ipv6 enable
switch (config)# interface vlan 50 ipv6 enable
```

Step 5. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::1 /64
switch (config)# interface vlan 30 ipv6 address 2103:db01::2 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::1 /64
```

Step 6. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2101:db01::2
```

Step 7. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2105:db01::2
```

➢ ***To configure Router2:***

Step 1. Disable prefix mode on the CLI. Run:

```
switch (config)# no cli default prefix-mode enable
```

Step 2. Enable the VLANs on the system. Run:

```
switch (config)# vlan 10
switch (config vlan 10) # exit
switch (config)# vlan 20
switch (config vlan 20) # exit
switch (config)# vlan 50
switch (config vlan 50) # exit
```

Step 3. Configure the switch ports to accept the VLANs of which they are part only. Run:

```
switch (config)#  1/1 switchport access vlan 10  // port2
switch (config)#  1/2 switchport access vlan 50  // port8
switch (config)#  1/36 switchport access vlan 20 // port5
```

Step 4. Disable spanning tree. Run:

```
switch (config)# no spanning-tree
```

Step 5. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

Step 6. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 20
switch (config interface vlan 20) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

Step 7. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::2 /64
switch (config)# interface vlan 20 ipv6 address 2102:db01::1 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::2 /64
```

Step 8. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2101:db01::1
```

Step 9. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2105:db01::1
```

➢ ***Ping neighbor to verify IPv6 configuration:***

```
switch (config)# ping6 2101:db01::2
PING 2101:db01::2(2101:db01::2) 56 data bytes
64 bytes from 2101:db01::2: icmp_seq=1 ttl=64 time=0.371 ms
64 bytes from 2101:db01::2: icmp_seq=2 ttl=64 time=0.620 ms
64 bytes from 2101:db01::2: icmp_seq=3 ttl=64 time=0.192 ms
64 bytes from 2101:db01::2: icmp_seq=4 ttl=64 time=0.277 ms
64 bytes from 2101:db01::2: icmp_seq=5 ttl=64 time=0.231 ms
```

6.2.3 Commands

ipv6 enable

	ipv6 enable **no ipv6 enable** Assigns automatic link-local IPv6 address to the interface. The no form of the command deassigns that automatic local address and disables IPv6 if no static IPv6 address has been assigned to the interface.
Syntax Description	N/A
Default	Unassigned
Configuration Mode	config interface vlan config interface loopback config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.4.1100 3.6.4110 Updated notes and command description.
Role	admin
Example	`switch (config vlan 10) # ipv6 enable`
Related Commands	
Note	Assigning an IPv6 address to an interface also enables IPv6 processing on the interface.

ipv6 address

ipv6 address <ipv6-address> /<length>
no ipv6 address [<ipv6-address> [/<length>]]

Enables IPv6 processing and assigns an IPv6 address to the interface.
The no form of the command removes the specified IPv6 address. If no address is specified, then all addresses of the interface are removed.

Syntax Description	ipv6-address	IPv6 address.
	length	Mask length for the associated address space. Range: 1-128. The mask length may be configured without a space (i.e. <ipv6-address>/<length>).
Default	N/A	
Configuration Mode	config interface vlan config interface loopback config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description and example output.
Role	admin	
Example	`switch (config vlan 10) # ipv6 address 2001::1 /120` `switch (config vlan 10) # ipv6 address 2001::1/120`	
Related Commands		
Note	An interface may have up to 16 IPv6 address assignments	

ipv6 nd managed-config-flag

	ipv6 nd managed-config-flag **no ipv6 nd managed-config-flag** Sets the managed address configuration flag in IPv6 router advertisements. The no form of the command restores the default setting.
Syntax Description	N/A
Default	Managed address configuration flag is not set
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.4.1100 3.6.4110 Updated configuration mode.
Role	admin
Example	`switch (config vlan 10) # ipv6 nd managed-config-flag`
Related Commands	
Note	

ipv6 nd ns-interval

ipv6 nd ns-interval <period>
no ipv6 nd ns-interval

Configures the interval between IPv6 neighbor solicitation (NS) transmissions. The no form of the command restores the default value.

Syntax Description	period	In milliseconds. Range: 1000-4294967295.
Default	1000 milliseconds	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ns-interval 1500`	
Related Commands		
Note		

ipv6 nd other-config-flag

	ipv6 nd other-config-flag **no ipv6 nd other-config-flag** Indicates that other configuration information is available via DHCPv6. The no form of the command removes the other configuration flag.
Syntax Description	N/A
Default	Not set
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.4.1100 3.6.4110 Updated configuration mode.
Role	admin
Example	`switch (config vlan 10) # ipv6 nd other-config-flag`
Related Commands	
Note	

ipv6 nd prefix

ipv6 nd prefix <ipv6-address> /<length> [no-autoconfig] [no-onlink] [valid-time {<time> | infinite}] [preferred-time {<time> | infinite}]
ipv6 nd prefix <prefix> no-advertise
no ipv6 nd prefix <prefix>

Configures inclusion for router advertisements (RAs) for neighbor.
The no form of the command removes the corresponding IPv6 nd prefix.

Syntax Description	ipv6-address	IPv6 address.
	length	Prefix length for the associated address space. Range: 1-128.
	no-advertise	Prevents advertising of the specified default prefix.
	valid-time	Time in seconds. Range: 0-4294967295.
	preferred-time	Time in seconds. Range: 0-4294967295.
	no-autoconfig	Indicates that the prefix cannot be used for stateless address configuration.
	no-onlink	Indicates that the prefix cannot be used for on-link determination
Default	valid-time: 2592000 seconds preferred-time: 604800 seconds no-autoconfig: Reset, autoconfig enabled no-onlink: Reset, on-link determination is enabled	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description, configuration mode and default values.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd prefix 2001::1 /120`	
Related Commands		
Note	• Valid time must be larger than preferred time • By default, the router advertises all configured sunbnets on the interface	

ipv6 nd ra dns-servers lifetime

ipv6 nd ra dns-servers lifetime {<time> | infinite}
no ipv6 nd ra dns-servers lifetime

Advertises a lifetime of a Recursive DNS Server (RDNSS).
Using RDNSS and DNSSL options, an IPv6 host can perform IPv6 address network configuration and DNS information simultaneously, without using DHCPv6 for the DNS configuration.
The no form of the command resets the lifetime value to default.

Syntax Description	time	Possible values: • 0 - RDNSS address can no longer be used • 1-4294967295 (sec)
	infinite	A value of all one bits (0xffffffff) and "infinite" represents infinity.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”.	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command and syntax description, configuration mode and default values.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra dns-servers lifetime infinite`	
Related Commands		
Note	• A lifetime value set for an individual RDNSS overrides this value. • The lifetime value is the maximum amount of time after a route advertisement packet is sent that the RDNSS referenced in the packet may be used for name resolution.	

ipv6 nd ra dns-server

ipv6 nd ra dns-server <ipv6 address> [lifetime [<time> | infinite]]
no ipv6 nd ra dns-server [<ipv6 address>]

Configures the IPv6 address of a Recursive DNS Server (RDNSS) to include in the neighbor-discovery router advertisements (RAs).
The no form of the command removes the RDNSS from the configuration.

Syntax Description	ipv6 address	IPv6 address of RDNSS
	lifetime	Maximum lifetime value for the specified RDNSS entry. Possible values: • 0 - RDNSS address can no longer be used • 1-4294967295 in seconds
	infinite	A value of all one bits (0xffffffff) and "infinite" represents infinity.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”.	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command, example and syntax description, configuration mode and default values.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra dns-server 2001::1 lifetime infinite`	
Related Commands		
Note	• Including RDNSS information in RAs provides DNS server configuration for connected IPv6 hosts without requiring DHCPv6 • Multiple servers can be configured on the interface by using the command repeatedly • A lifetime value for the RDNSS can optionally be specified with this command, and overrides any default value configured for the interface using the ipv6 nd ra dns-servers lifetime command	

ipv6 nd ra dns-suffixes lifetime

ipv6 nd ra dns-suffixes lifetime {<time> | infinite}
no ipv6 nd ra dns-suffixes lifetime

Advertises a lifetime of a DNS Search List (DNSSL).
Using RDNSS and DNSSL options, an IPv6 host can perform IPv6 address network configuration and DNS information simultaneously, without using DHCPv6 for the DNS configuration.
The no form of the command resets the lifetime value to its default.

Syntax Description	time	Possible values: • 0 – RDNSS address can no longer be used • 1-4294967295 – in seconds
	infinite	A value of all one bits (0xffffffff) and "infinite" represents infinity.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”.	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command, example and syntax description, configuration mode and default values.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra dns-suffix mellanox.com lifetime infinite`	
Related Commands		
Note	• The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries	

ipv6 nd ra dns-suffix

ipv6 nd ra dns-suffix <domain-name> [lifetime {<time> | infinite}]
no ipv6 nd ra dns-suffix [<domain-name>]

Creates a DNS search list (DNSSL) to include in the neighbor-discovery Router Advertisements (RAs).
The no form of the command removes the DNSSL from the configuration.

Syntax Description	domain-name	Domain suffix for IPv6 hosts to append to short unqualified domain names for DNS queries. The suffix must contain only alphanumeric characters, "." (periods), "-" (hyphens), and must begin and end with an alphanumeric character.
	lifetime	Maximum lifetime value for the specified DNSSL entry.
	time	Possible values: • 0 – DNSSL must not be used for name resolution • 1-4294967295 – in seconds
	infinite	A value of all one bits (0xffffffff) and "infinite" represents infinity.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command "ipv6 nd ra interval".	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command, example and syntax description, configuration mode and default values.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra dns-suffix mellanox.com lifetime infinite`	
Related Commands		
Note	• The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries • Multiple DNS domain names can be added to the DNSSL by reusing the command • A lifetime value for the DNSSL can optionally be specified with this command which overrides any default value configured for the interface using the command "ipv6 nd ra dns-suffixes lifetime"	

ipv6 nd ra hop-limit

ipv6 nd ra hop-limit <limit>
no ipv6 nd ra hop-limit

Sets a suggested hop-limit value to be included in route advertisement (RA) packets. The no form of the command resets the parameter to its default value.

Syntax Description	limit	The hop-limit value to be included by attached hosts in outgoing packets. • 0 – unspecified (by this router) • 1-255 – number of hops
Default	Limit value is 64	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration modes.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra hop-limit 70`	
Related Commands		
Note		

ipv6 nd ra interval max-period

ipv6 nd ra interval max-period <time> [min-period <time>]
no ipv6 nd ra interval

Configures the interval between IPv6 router advertisement (RA) transmissions. The no form of the command resets the parameter to its default value.

Syntax Description	time	Maximum interval between successive IPv6 router advertisement transmissions. Range: 4-1800 seconds.
	min-period	minimum interval between successive IPv6 router advertisement transmissions. • No parameter: Default is used • 4-1800
Default	max-period: 600 seconds min-period: See Note	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description, configuration modes and notes.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra interval max-period 600`	
Related Commands		
Note	• The min-period must be 0.33 * <max-period> if <max-period> is >= 9 seconds; otherwise, the default is Router Advertisement Interval • The parameter min-period must be no less than 3 seconds and no greater than 0.75*max-period	

ipv6 nd ra lifetime

ipv6 nd ra lifetime <time>
no ipv6 nd ra lifetime

Router lifetime is associated with a router's usefulness as default route, it does not apply to information contained in other message fields or options. Options that need time limits for their information include their own lifetime fields.
The no form of the command resets the parameter to its default value.

Syntax Description	time	The router lifetime specifies the period that the router can be considered as a default router by RA recipients in seconds. • 0 – the router should not be considered a default router on this interface • 1-9000 – lifetime period advertised in RAs should not be less than the max router advertisement interval
Default	3*<router advertisement interval>	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Added support for IPv6
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra lifetime 300`	
Related Commands		
Note		

ipv6 nd ra mtu suppress

	ipv6 nd ra mtu suppress **no ipv6 nd ra mtu suppress** Suppresses advertisement (RA) MTU option sent to router. MTU option ensures all nodes on a link use the same MTU value. The no form of the command restores the MTU option to enabled.
Syntax Description	N/A
Default	Suppressed
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.4.1100 3.6.4110 Updated command Syntax and configuration mode.
Role	admin
Example	`switch (config vlan 10) # ipv6 nd ra mtu suppress`
Related Commands	
Note	If not suppressed, MTU of the interface is advertised.

ipv6 nd ra suppress

	ipv6 nd ra suppress [all] **no ipv6 nd ra suppress** Suppresses periodic and solicited IPv6 router advertisement (RA) transmissions. The no form of the command restores the transmission of RAs.	
Syntax Description	all	Configures the switch to suppress all RAs, including those responding to a router solicitation.
Default	Only unsolicited RAs transmitted periodically are suppressed	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax and configuration mode.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra suppress all`	
Related Commands		
Note		

ipv6 nd reachable-time

	ipv6 nd reachable-time <time> **no ipv6 nd reachable-time** Sets the time period the switch includes in the reachable time field of out-going advertisements (RAs). The no form of the command resets the parameter to its default value.	
Syntax Description	time	In milliseconds; the reachable time defines the period that a node assumes a neighbor is reachable after having received a reachability confirmation. Values: • 0 - unspecified by router • 1 - 3600000 the period that a node assumes a neighbor is reachable.
Default	0 (unspecified)	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax, configuration mode and notes.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd reachable-time 30000`	
Related Commands		
Note	• RAs that advertise zero seconds indicate that the router does not specify a reachable time	

ipv6 nd router-preference

	ipv6 nd router-preference {high | medium | low} **no ipv6 nd router-preference** Sets the value the switch enters in the default router preference (DRP) field of router advertisements (RAs) it sends. The no form of the command resets the parameter to its default value.
Syntax Description	N/A
Default	Medium
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.4.1100 3.6.4110 Updated configuration modes.
Role	admin
Example	`switch (config vlan 10) # ipv6 nd router-preference high`
Related Commands	
Note	• IPv6 hosts maintain a default router list from which to select a router for traffic to offlink destinations. The router's address is then saved in the destination cache. The neighbor discovery protocol (NDP) prefers routers that are reachable or probably reachable over routers whose reachability is unknown or suspect. For reachable or probably reachable routers, NDP can either select the same router every time or cycle through the router list. DRP values specify a host's preferred router. • If router lifetime is zero, preference value must be medium

ipv6 nd retrans-timer

	ipv6 nd retrans-timer <time> **no ipv6 nd retrans-timer** Advertises the time between consecutive neighbor solicitation (NS) messages. The no form of the command resets the parameter to its default value.	
Syntax Description	time	In milliseconds; the time between retransmitted neighbor solicitation messages. Possible values: • 0 – unspecified • Range – 1000-4294967295
Default	0 (unspecified)	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax, configuration mode and example output.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd retrans-timer 1000`	
Related Commands		
Note		

ipv6 nd redirects

	ipv6 nd redirects **no ipv6 nd redirects** Enables sending ICMPv6 redirect messages. The no form of the command disables sending ICMPv6 redirect messages.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan
History	3.4.1100
Role	admin
Example	`switch (config interface vlan 10) # ipv6 nd redirects`
Related Commands	
Note	

ipv6 nd dad attempts

ipv6 nd dad attempts <number>
no ipv6 nd dad attempts

Sets the number of consecutive neighbor solicitation messages sent for duplicate address detection (DAD) validation.
The no form of the command resets the value to its default.

Syntax Description	number	Number of attempts: • 0 – DAD is not performed • Valid range: 1-1000
Default	1	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode.
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd dad attempts 10`	
Related Commands		
Note		

ipv6 neighbor

ipv6 neighbor [vrf <name>] <ipv6-address> <mac-address>
ipv6 neighbor <ipv6-address> interface {ethernet <port> | vlan <vlan-id> | port-channel <port-channel>} <mac-address>
no interface {ethernet <port> | vlan <vlan-id> | port-channel} ipv6 neighbor <ipv6-address> <mac-address>
no ipv6 neighbor [vrf <name>] <ipv6-address>

Creates an IPv6 neighbor discovery cache static entry.
The no form of the command removes the specified static entry from the IPv6 neighbor discovery cache.

Syntax Description	ipv6-address	IPv6 address
	ethernet <port>	Ethernet port. Format <slot>/<port>.
	vlan <vlan-id>	VLAN ID
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
	3.6.4110	Updated command syntax.
Role	admin	
Example	`switch (config vlan 10) # ipv6 neighbor 2001:db01::1 vlan 10 4:4:4:4:4:4`	
Related Commands		
Note	This command do not affect any dynamic entries in the cache.	

clear ipv6 neighbors

clear ipv6 neighbors {ethernet <slot> /<port> | port-channel <port-channel> | vlan <vlan-id>} [<ipv6-addr>]

Removes the specified dynamic IPv6 neighbor discovery cache entries.

Syntax Description	ethernet	Ethernet port. Format: <slot>/<port>.
	vlan	VLAN interface
	ipv6-addr	IPv6 address
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
	3.6.4110	Updated command
Role	admin	
Example	`switch (config) # clear ipv6 neighbors ethernet 1/4`	
Related Commands		
Note	• Commands that do not specify an IPv6 address remove all dynamic entries for the listed interface • Commands that do not specify an interface remove all dynamic entries	

ipv6 route

- General route:
  ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} <next-hop-ipv6-address> [<distance>]
- Local route:
  ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} {<ethernet <port> | vlan <id> | port-channel <id>} [<distance>]
- Drop route:
  ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} null0 [<distance>]
- Delete route(s):
  no ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} [<next-hop-ipv6-address>]

Creates an IPv6 static route.
The no form of the command deletes static routes.

Syntax Description	ipv6-address	IPv6 address.
	ipv6-prefix	IPv6 address + mask length without space, e.g. a1:a2::33/64.
	length	Prefix length for the associated address space. Range: 1-128.
	next-hop-ipv6-address	IPv6 address of the next-hop
	distance	Administrative distance assigned to route. Options include: • No parameter – route is assigned a default administrative distance of 1 • 1-255 – the administrative distance assigned to route
	null0	Creates a black hole route with action DROP
Default	No distance parameter indicated: Administrative distance of 1	
Configuration Mode	config	
History	3.4.1100	
	3.6.4110	Updated command
Role	admin	
Example	`switch (config) # ipv6 route 3003:db01:: /64 2001:db01::1` `switch (config) #`	

Related Commands	
Note	• Static routes have a default administrative distance of 1 • Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data. • Multiple routes which are configured to the same destination with the same administrative distance comprise an Equal Cost Multi-Path (ECMP) route • A no command not including a source deletes all statements to the destination • Route with distance value 255 is not inserted to the forwarding table

ipv6 routing

	ipv6 routing **no ipv6 routing** Enables forwarding IPv6 unicast packets. The no form of the command disables IPv6 unicast routing.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.1100
Role	admin
Example	`switch (config) # ipv6 routing`
Related Commands	
Note	• When routing is enabled, the switch attempts to deliver inbound packets to destination addresses by forwarding them to interfaces or next hop addresses specified by the IPv6 routing table

show ipv6 interfaces

show ipv6 interfaces [{{ethernet <port> | port-channel <port-channel> | vlan <vlan-id>}}| brief]

Displays the status of specified routed interfaces that are configured for IPv6.

Syntax Description	ethernet <port>	Displays output pertaining to the specified Ethernet interface
	port-channel <port-channel>	Displays output pertaining to the specified LAG interface
	vlan <vlan-id>	Displays output pertaining to the specified VLAN interface
	brief	Shows basic IPv6 information regarding all IPv6 interfaces
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Role	admin	

Example

```
switch (config) # show ipv6 interface

Vlan10 is Enabled ,  line protocol is UP
        IPv6 : Enabled
        Link-local address : fe80::f652:14ff:fe2d:9808
        Global Unicast Addresses :
        2001:db01::2  /64
        Joined Group Addresses :
        ff02::1
        ff02::2
        ff02::1:ff2d:9808
        MTU : 1500 bytes
        ICMP error messages limited to every milliseconds : 100
        ICMP redirects : enabled
        ND DAD : enabled
        Number of DAD attempts : 1
        ND reachable time (milliseconds) : 30000
        ND advertised retransmit interval (milliseconds) : 0
        ND router advertisements maximum interval (seconds) : 600
        ND router advertisements minimum interval (seconds) : 198
        ND router advertisements managed configuration flag : unset
        ND router advertisements other   configuration flag : unset
        ND solicited router advertisement : suppressed
        ND router advertisements lifetime (seconds) : 1800
        ND advertised default router preference : medium
        ND router advertisements hop-limit : 64
switch (config) #
```

Related Commands
Note

show ipv6 interfaces brief

show ipv6 interfaces [<type> <id>] brief

Displays basic IPv6 information regarding all IPv6 interfaces

Syntax Description	<type> <id>	Specifies the interface for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ipv6 interface brief

-------------------------------------------------------------------------------------------------------------
Interface   Address/Mask       Primary          Address-state   Admin-state    Oper-state    MTU     VRF
-------------------------------------------------------------------------------------------------------------
mgmt0       fe80::784e/64                       valid           Enabled        Up            1500    default
Eth1/1      2001::1/64         primary          valid           Enabled        Down          1500    default
Eth1/1      2002::1/64                          valid
```

Related Commands

Note

show interfaces null0

show interfaces null0 [vrf <vrf-name>]

Displays blackhole route byte and packet counters.

Syntax Description	N/A N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Role	admin

Example	`switch (config) # show interfaces null0`
	`10 packets`
	`740 bytes`
	`switch (config) #`

Related Commands	
Note	

show ipv6 neighbors

show ipv6 neighbors [{ethernet <port> | port-channel <port-channel> | vlan <vlan-id>} | <ipv6 address> | summary]

Displays IPv6 neighbor discovery (ND) cache information.

Syntax Description	ethernet <port>	Shows output pertaining to the specified Ethernet interface.
	vlan <vlan-id>	Shows output pertaining to the specified VLAN interface.
	ipv6 address	IPv6 address of individual neighbor
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.4110	Updated command syntax and example output.
Role	admin	

Example	`switch (config) # show ipv6 neighbors`
	`IPv6 Address MAC Address State Interf`
	`---------- ---------- ---------- ------`
	`2001:db01::1 f4:52:14:2d:98:88 Reachable vlan10`
	`switch (config) #`	

Related Commands		
Note		

show ipv6 route

show ipv6 route [vrf <vrf-name] {[<ipv6-address> <ipv6-address>/<length> [longer-prefixes]][connected | bgp | static] }

Displays IPv6 neighbor discovery (ND) cache information.

Syntax Description	ipv6-addr	Filters routes by IPv6 address or prefix
	longer-prefixes	Displays output for longer prefix entries
	connected	Displays entries for routes to networks directly connected to the switch
	static	Displays entries added through CLI commands
	summary	Displays the current contents of the IPv6 routing table in summary format
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.4110	Updated Example
	3.6.8008	Updated Example
Role	admin	
Example		
Related Commands		
Note		

```
switch (config) # show ipv6 route

Flags:
  F: Failed to install in H/W
  B: BFD protected
  i: BFD session initializing
  x: protecting BFD session failed

VRF Name default:
  ------------------------------------------------------------------------
  Destination        Flag   Gateway           Interface      Source   AD/M
  ------------------------------------------------------------------------
  fe80::/64                 ::                mgmt0          direct   256/256
  default                   ::                mgmt0          direct   1/1
```

6.3 OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol for IP networks. It uses a link state routing algorithm and falls into the group of interior routing protocols, operating within a single autonomous system (AS).

OSPF-speaking routers send Hello packets on all OSPF-enabled IP interfaces. If two routers sharing a common data link agree on certain parameters specified in their respective Hello packets, they become neighbors.

Adjacencies, which can be thought of as virtual point-to-point links, are formed between some neighbors. OSPF defines several network types and several router types. The establishment of an adjacency is determined by the types of routers exchanging Hellos and the type of network over which the Hello packets are exchanged.

Each router sends link-state advertisements (LSAs) over all adjacencies. The LSAs describe all of the router's links, or interfaces, the router's neighbors, and the state of the links. These links might be to stub networks (those without another router attached), to other OSPF routers, to networks in other areas, or to external networks (those learned from another routing process). Because of the varying types of link-state information, OSPF defines multiple LSA types.

Each router receiving an LSA from a neighbor records the LSA in its link-state database and sends a copy of the LSA to all of its other neighbors. By flooding LSAs throughout an area, all routers will build identical link-state databases.

When the databases are complete, each router uses the SPF algorithm to calculate a loop-free graph describing the shortest (lowest cost) path to every known destination, with itself as the root.

When all link-state information has been flooded to all routers in an area, and neighbors have verified that their databases are identical, it means the link-state databases have been synchronized and the route tables have been built. Hello packets are exchanged between neighbors as keepalives, and LSAs are retransmitted. If the network topology is stable, no other activity should occur.

For OSPF network design over Mellanox L2 VMS, please refer to Mellanox Virtual Modular Switch Reference Guide.

6.3.1 Router ID

The router ID is a 32-bit number assigned to the router running the OSPF protocol. This number uniquely identifies the router in the OSPF link-state database.

Router ID can be configured statically, however, if it is not configured, then the default election is as follows:

- If a loopback interface already exists, the router ID takes the loopback IP address;
- Otherwise, the lowest IP address is elected as router ID

6.3.2 ECMP

Equal-cost multi-path (ECMP) routing is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple paths. The OSPF link-state routing algorithm can

find multiple routes to the same destination, all multiple routes are added to the routing table only if those routes are equal-cost routes.

In case there are several routes with different cost, only the route with the lowest cost is selected. In case there are multiple routes with the same lowest cost, all of them are used (up to maximum of 64 ECMP routes).

ECMP is not configurable but is enabled by default for OSPF.

6.3.3 Configuring OSPF

Figure 39: OSPF Basic Topology

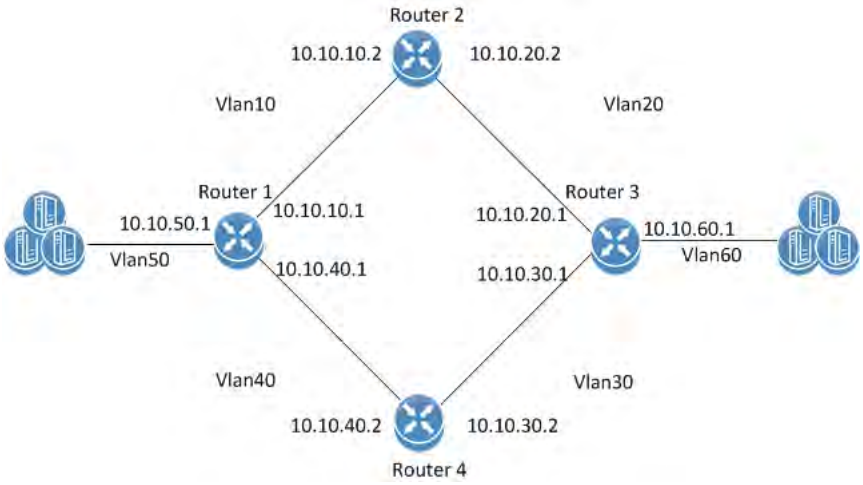

Precondition steps:

The following configuration example refers to Router 2 in Figure 39. The remainder of the routers in the figure are configured similarly.

It is recommended to disable STP before enabling OSPF. Use the command `no spanning-tree`.

Step 1. Enable IP routing functionality. Run:.

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:.

```
switch (config)# vlan 10
switch (config)# vlan 20
```

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config ethernet 1/1)# switchport access vlan 10
switch (config ethernet 1/1)# exit
switch (config)# interface ethernet 1/2
switch (config ethernet 1/2)# switchport access vlan 20
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

Step 5. Apply IP address to the VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /16
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

Step 7. Create a second VLAN interface. Run:

```
switch (config)# interface vlan 20
```

Step 8. Apply IP address to the second VLAN interface. Run:

```
switch (config interface vlan 20)# ip address 10.10.20.2 /16
```

Step 9. Enable the second interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

Basic OSPF Configuration:

Step 1. To enable OSPF configuration run:

```
switch (config)# protocol ospf
```

Step 2. To create a router OSPF instance run:

```
switch (config)# router ospf
```

Only one instance of OSPF per VRF is supported.

Step 3. Associate the VLAN interfaces to the OSPF area. Area 0 is the backbone area, run:

```
switch (config interface vlan 10)# ip ospf area 0
switch (config interface vlan 10)# exit
switch (config)# interface vlan 20
switch (config interface vlan 20)# ip ospf area 0
```

6.3.4 Verifying OSPF

➢ ***To verify OSPF configuration and status:***

Step 1. Verify OSPF configuration and status. Run:

```
switch (config) # show ip ospf

Routing Process 1 with ID 10.10.10.10 vrf-default

Stateful High Availability disabled
Graceful-restart is not supported
Supports only single TOS (TOS 0) route
Opaque LSA not supported
OSPF Admin State is enabled
Redistributing External Routes: Disabled
Administrative distance 110
Reference Bandwidth is 40Gb
Initial SPF schedule delay 1 msecs
SPF Hold time 10 msecs
Maximum paths to destination 64
Router is not originating router LSA with maximum metric
Condition: Always
Number of external LSAs  0, checksum sum  0
Number of opaque AS LSAs 0,checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 nssa
Number of active areas is 1, 1 normal, 0 stub, 0 nssa

Area (0.0.0.0) (Active)
Interfaces in this area: 2 Active Interfaces: 2
Passive Interfaces: 0
SPF Calculation has run 5 times
This area is Normal area
Number of LSAs: 1, checksum sum 7700

switch (config) #
```

Step 2. Verify the OSPF neighbors status. Make sure that each neighbor reaches FULL state with its peer to enable it take part in all dynamic routing changes in the network. Run:

```
switch (config) # show ip ospf neighbors

Neighbor 10.10.10.1, interface address 10.10.10.2
In the area 0.0.0.0 via interface Vlan 10
Neighbor priority is 1, State is FULL
BDR is 10.10.10.1
Options 0
Dead timer due in 35
```

```
Neighbor 10.10.20.1, interface address 10.10.20.2
In the area 0.0.0.0 via interface Vlan 20
Neighbor priority is 1, State is FULL
BDR is 10.10.20.1
Options 0
Dead timer due in 35

switch (config) #
```

Step 3. Verify the OSPF Interface configuration and status run:

```
switch (config) # show ip ospf interface

Interface Vlan is 10 Enabled, line protocol is Down
IP address 10.10.10.2, Mask 255.255.0.0 [primary]
Process ID 1 VRF Default, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DOWN, Network Type BROADCAST, Cost 1
Transmit delay 1 sec, Router Priority 1
No designated router on this network
No backup designated router on this network
Timer intervals (sec's): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0

Interface Vlan is 20 Enabled, line protocol is Up
IP address 10.10.20.2, Mask 255.255.0.0 [primary]
Process ID 1 VRF Default, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DESIGNATED ROUTER, Network Type BROADCAST, Cost 1
Transmit delay 1 sec, Router Priority 1
No designated router on this network
No backup designated router on this network
Timer intervals (sec's): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0

switch (config) #
```

6.3.5 Commands

6.3.5.1 Config

protocol ospf

	protocol ospf **no protocol ospf** Enables Open Shortest Path First Protocol (OSPF), and unhides the related OSPF commands. The no form of the command deletes the OSPF configuration and hides the OSPF related commands.
Syntax Description	N/A
Default	OSPF feature is disabled.
Configuration Mode	config
History	3.3.3500
Role	admin
Example	`switch (config)# protocol ospf`
Related Commands	ip routing
Note	

router ospf

router ospf [<process-id> [vrf <vrf-name>]]
no router ospf [<process-id> [vrf <vrf-name>]]

Enters router OSPF configuration mode, and creates default OSPF instance on specific VRF with specific Process ID if one does not exist.
The no form of the command deletes the OSPF instance.

Syntax Description	process-id	OSPF instance ID
	vrf	VRF name (e.g. default)
Default	Process ID: 1 VRF: Active VRF routing-context	
Configuration Mode	config	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Role	admin	
Example	`switch (config)# router ospf 2 vrf myvrf` `switch (config router ospf 2)#`	
Related Commands	N/A	
Note	Only one OSPF instance per VRF is supported.	

6.3.5.2 Config Router

router-id

router-id <ip-address>
no router-id

Sets Router ID for the OSPF instance.
The no form of the command causes automatic election of router ID by the router.

Syntax Description	ip-address	The Router id in IP address format.
Default	The router ID is a 32-bit number assigned to the router running the OSPF protocol. This number uniquely identifies the router within an OSPF link-state database. Router ID can be configured statically, however, if it is not configured, then the default election is as follows: • If a loopback interface already exists, the router ID takes the loopback IP address; • Otherwise, the lowest IP address is elected as router ID.	
Configuration Mode	config ospf router	
History	3.3.3500	
	3.7.11xx	Updated default
Role	admin	
Example	`switch (config router ospf)# router-id 10.10.10.10`	
Related Commands	N/A	
Note		

shutdown

	shutdown **no shutdown** Disables the OSPF instance. The no form of the command enables the OSPF instance.
Syntax Description	N/A
Default	Enable (no shutdown)
Configuration Mode	config ospf router
History	3.3.3500
Role	admin
Example	`switch (config router ospf)# shutdown`
Related Commands	N/A
Note	

auto-cost reference-bandwidth

auto-cost reference-bandwidth <ref-bw> [Gbps | Mbps]
no auto-cost reference-bandwidth

Configures reference-bandwidth in Gb/s (Default) or Mb/s.
The no form of the command resets this parameter to its default value.

Syntax Description	ref-bw	Range: 1-4294
	Gbps	Value in Gb/s (default if not specified)
	Mbps	Value in Mb/s
Default	100 Gb/s	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# auto-cost reference-bandwidth 10 Gbps`	
Related Commands	N/A	
Note		

distance

distance <value>
no distance

Configures the OSPF route administrative distance.
The no form of the command resets this parameter to default.

Syntax Description	value	OSPF administrative distance. Range is 1-255.
Default	110	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# distance 100`	
Related Commands	N/A	
Note		

redistribute

redistribute {bgp | direct | static | ebgp | ibgp}
no redistribute {bgp | direct | static}

Enables importing routes from other routing protocols as well as any statically configured routers into OSPF.
The no form of the command disables the importing of the routes.

Syntax Description	direct	Redistribute directly connected routes
	bgp	Redistribute routes from BGP protocol
	ibgp	Redistribute IBGP routes
	ebgp	Redistribute EBGP routes
	static	Redistribute static configured routes
Default	Disable (no redistribution)	
Configuration Mode	config ospf router	
History	3.6.3506 3.2.1000	
Role	admin	
Example	`switch (config router ospf)# redistribute direct`	
Related Commands	N/A	
Note	Routes from multiple protocols can be imported in parallel.	

timers throttle spf

	timers throttle spf <spf-delay> <spf-hold> **no timers throttle spf** Sets the OSPF throttle SPF timers. The no form of the command resets the timers to default.	
Syntax Description	spf-delay	The interval by which SPF calculations delayed after a topology change reception. Range is 0-100 milliseconds.
	spf-hold	The minimum delay between two consecutive delay calculations. Range is 0-1000 milliseconds.
Default	spf-delay: 1 millisecond spf-hold: 10 millisecond	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# timers throttle spf 100 1000`	
Related Commands	N/A	
Note		

area default-cost

area <area-id> default-cost <cost>
no area <area-id> default-cost

Specifies cost for the default summary route sent into an OSPF stub or not-so-stubby area (NSSA).
The no form of the command sets the cost to the default value.

Syntax Description	area-id	OSPF area-id. Range is 0-4294967295.
	cost	The cost for the default summary route. Range is 1-16777215.
Default	The summary route cost is based on the area border router that generated the summary route.	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# area 0 default-cost 100`	
Related Commands	N/A	
Note	Base cost for all calculation is 56GbE.	

area range

area <area-id> range <ip-address> <prefix> [not-advertise]
no area <area-id> range <ip-address> <prefix> [not-advertise]

Consolidates and summarizes routes at an OSPF area boundary.
The no form of the command removes the ip-prefix range from summarization.

Syntax Description	area-id	OSPF area-ID. Range is 0-4294967295.
	ip-address	IP Address.
	not-advertise	Suppresses routes that match the specified IP address.
	prefix	Network prefix (in the format of /24, or 255.255.255.0 for example).
Default	Disabled	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# area 0 range 10.10.10.10 /24`	
Related Commands	N/A	
Note		

area stub

area <area-id> stub [no-summary]
no area <area-id> stub [no-summary]

Configures an area as an OSPF stub area (an area is created if non-existent).
The no form of the command removes the stub area configuration and changes the area to normal, or deletes the area (if stub is not used).

Syntax Description	area-id	OSPF area-ID. Range is 0-4294967295.
	no-summary	Summary route will not be advertised into the stub area.
Default	Summary route will be advertised.	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# area 0 stub`	
Related Commands	N/A	
Note		

area nssa

area <area-id> nssa [default-information-originate [metric <m-value>] [metric-type <m-type>]] [nosummary] [translate type7 always]
no area <area-id> nssa [default-information-originate] [no-summary] [translate type7 always]

Configures an area as an OSPF not-so-stubby (NSSA) area.
The no form of the command removes the NSSA area configuration and changes the area to default.

Syntax Description	area-id	OSPF area ID. Range is 0-4294967295.
	default-information-originate	A default type7 LSA (Link State Advertisements) is generated into the NSSA area.
	m-type	Metric type for OSPF. Range is 1-2.
	m-value	Metric value for OSPF. Range is 1-65535.
	no-summary	Summary route will not be advertised into the NSSA area.
	translate type7 always	Type7 LSAs is translated to type5 LSAs (Link State Advertisements).
Default	Default m-type:2 Default m-value:10	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# area 0 nssa`	
Related Commands	N/A	
Note	An area can be either stub, NSSA or normal.	

no area

no area <area-id>

Deletes OSPF area and its related configuration.

Syntax Description	area-id	OSPF area ID Range is 0-4294967295
Default	N/A	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# no area 1`	
Related Commands	N/A	
Note	The command fails if the area is attached to active interfaces.	

default-information originate

default-information originate [always] [metric <m-value>] [metric-type <m-type>]
no default-information originate

Enables default route origination to normal areas.
The no form of the command resets the parameter values to their default.

Syntax Description	always	Default route is always advertised even if the default route is not in the routing table
	metric	Route metric value. Range: 1-65535.
	metric-type	Metric type. Range: 1-2.
Default	m-value – 1 m-type – 2	
Configuration Mode	config ospf router	
History	3.6.8008	
Role	admin	
Example	`switch (config router ospf)# default-information originate always`	
Related Commands	N/A	
Note	When default route origination is enabled, the router automatically becomes ASBR and advertises a default route	

summary-address

summary-address <ip-address> <prefix> [not-advertise]
no summary-address <ip-address> <prefix> [not-advertise]

Creates aggregate addresses for the OSPF protocol.
The no form of the command disables the aggregation of the ip-address.

Syntax Description	ip-address	The summary IP address.
	not-advertise	Suppresses routes that match the specified ip-address.
	prefix	Network prefix (in the format of /24 or 255.255.255.0, for example).
Default	N/A	
Configuration Mode	config ospf router	
History	3.3.3500	
Role	admin	
Example	`switch (config router ospf)# summary-address 10.10.10.10 /24`	
Related Commands	N/A	
Note	Maximum of 1500 summarized IP addresses can be configured.	

6.3.5.3 Interface

ip ospf cost

ip ospf cost <cost>
no ip ospf cost

Sets OSPF cost of sending packet of this interface.
The no form of the command resets this parameter to default.

Syntax Description	cost	The Interface cost used by the OSPF. Range is 1-65535.
Default	Reference_BW/Link_BW	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
	3.7.11xx	Updated default
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf cost 100`	
Related Commands	N/A	
Note		

ip ospf dead-interval

ip ospf dead-interval <seconds>
no ip ospf dead-interval

Configures the interval during which at least one Hello packet must be received from a neighbor before the router declares that neighbor as down.
The no form of the command resets this parameter to its default.

Syntax Description	seconds	The dead-interval timer, in seconds. Range is 1-65535.
Default	40	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf dead-interval 10`	
Related Commands	N/A	
Note	The value must be the same for all nodes on the network.	

ip ospf hello-interval

ip ospf hello-interval <seconds>
no ip ospf hello-interval

Configures the interval between Hello packets that OSPF sends on the interface. The no form of the command resets this parameter to default.

Syntax Description	seconds	The Hello interval timer, in seconds. Range is 1-65535.
Default	10	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf hello-interval 20`	
Related Commands	N/A	
Note	The value must be the same for all nodes on the network.	

ip ospf priority

ip ospf priority <number>
no ip ospf priority

Configures the priority for this OSPF interface.
The no form of the command resets this parameter to default.

Syntax Description	number	The Interface priority used by the OSPF protocol. Range is 0-255
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf priority 100`	
Related Commands	N/A	
Note	• Use the “ip ospf priority” command to set the router priority, which determines the designated router for this network. When two routers are attached to a network, both attempt to become the designated router. • The router with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero cannot become the designated router or backup designated router.	

ip ospf network

ip ospf network <type>
no ip ospf network

Sets the OSPF interface network type.
The no form of the command resets the interface network type to its default.

Syntax Description	type	The network type on this interface. • broadcast • point-to-point
Default	broadcast for VLAN interfaces point-to-point for router port interfaces	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf network point-to-point`	
Related Commands	N/A	
Note	• The network type influences the behavior of the OSPF interface. An OSPF network type is usually broadcast, which uses OSPF multicasting capabilities. Under this network type, a designated router and backup designated router are elected. For point-to-point networks, there are only two neighbors and multicast is not required. • All routers on the same network must have the same network type.	

ip ospf retransmit-interval

ip ospf retransmit-interval <seconds>
no ip ospf retransmit-interval

Configures the time between OSPF link-state advertisement (LSA) retransmissions for adjacencies that belongs to the interface.
The no form of the command resets this parameter to its default.

Syntax Description	seconds	The retransmit interval in seconds. Range is 0-3600.
Default	5	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf retransmit-interval 10`	
Related Commands	N/A	
Note		

ip ospf passive-interface

	ip ospf passive-interface **no ip ospf passive-interface** Suppresses flooding of OSPF routing updates on an interface. The no form of the command reverts the status to active OSPF interface.
Syntax Description	N/A
Default	Active interface (no ip ospf passive-interface)
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Role	admin
Example	`switch (config interface vlan 10)# ip ospf passive-interface`
Related Commands	N/A
Note	

ip ospf transmit-delay

ip ospf transmit-delay <seconds>
no ip ospf transmit-delay

Sets the estimated time required to send an OSPF link-state update packet.
The no form of the command resets this parameter to its default.

Syntax Description	seconds	The transmit-delay interval in seconds. Range is 0-3600.
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf transmit-delay 2`	
Related Commands	N/A	
Note		

ip ospf shutdown

	ip ospf shutdown **no ip ospf shutdown** Disables the OSPF instance on the interface. The no form of the command enables the OSPF on this interface.
Syntax Description	N/A
Default	Enabled (no shutdown)
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Role	admin
Example	`switch (config interface vlan 10)# ip ospf shutdown`
Related Commands	N/A
Note	

ip ospf authentication

ip ospf authentication [message-digest]
no ip ospf authentication

Specifies the authentication type for OSPF.
The no form of the command disables the authentication.

Syntax Description	message-digest	Specifies that message-digest authentication (MD5) is used.
Default	Disabled (no)	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf authentication`	
Related Commands	N/A	
Note	• Without message-digest option, a simple password authentication will be used. • Message-digest authentication can be enabled only if a key is configured.	

ip ospf authentication-key

ip ospf authentication-key [<auth-type>] <password>
no ip ospf authentication-key

To assign a password for simple password authentication for the OSPF.
The no form of the command deletes the simple password authentication key.

Syntax Description	auth-type	The authentication type: 0 – unencrypted password 7 – MD5 key
	password	Authentication password (up to 8 alphanumeric string)
Default	Unencrypted password	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf authentication-key 0 mycleartextpassword`	
Related Commands	N/A	
Note	• When selecting an encrypted password “7”, the user must input a password encrypted with an MD5 key. • When selecting an unencrypted password “0”, the user must input a cleartext password. Then when examining the running-config, it exhibits the encrypted password.	

ip ospf message-digest-key

ip ospf message-digest-key <key-id> md5 [auth-type] <key>
no ip ospf message-digest-key <key-id>

Sets the message digest key for MD5 authentication.
The no form of the command deletes the key for MD5 authentication.

Syntax Description	auth-type	The authentication type: 0 - Unencrypted password 7 - MD5 key
	key	Authentication password, up to 8 alphanumeric string.
	key-id	Alphanumeric password of up to 16 bytes.
Default	Unencrypted (no)	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf message-digest-key mykeyid md5 7 mykey`	
Related Commands	N/A	
Note	The user cannot delete the last key until authentication is disabled.	

ip ospf area

ip ospf area <area-id>
no ip ospf area

Sets OSPF area of this interface (and creates the area if non-existent).
The no form of the command removes the interface from the area.

Syntax Description	area-id	OSPF area ID Range is 0-4294967295
Default	N/A	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface) config interface loopback	
History	3.3.3500	
Role	admin	
Example	`switch (config interface vlan 10)# ip ospf area 0`	
Related Commands	N/A	
Note		

6.3.5.4 Show

show ip ospf

show ip ospf [<process-id> [vrf <vrf-name>]]

Displays general OSPF configuration on specific VRF and status.

Syntax Description	process-id	OSPF instance ID
	vrf	VRF instance
Default	Process ID: 1 VRF: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Role	admin	

Example	`switch (config)# show ip ospf 2 vrf myvrf`
	
	`Routing Process 2 with ID 2.2.2.2 myvrf`
	
	`Stateful High Availability is not supported`
	`Graceful-restart is not supported`
	`Supports only single TOS (TOS 0) route`
	`Opaque LSA not supported`
	`OSPF Admin State is enabled`
	`Redistributing External Routes: Disabled`
	`Administrative distance 110`
	`Reference Bandwidth is 40 Gbps`
	`Initial SPF schedule delay 1 msecs`
	`SPF Hold time 5000 msecs`
	`Maximum paths to destination 64`
	`Router LSA with maximum metric is not supported`
	`Condition: Always`
	`Number of external LSAs 0, checksum sum 0`
	`Number of opaque AS LSAs 0, checksum sum 0`
	`Number of areas is 1, 1 normal, 0 stub, 0 nssa`
	`Number of active areas is 1, 1 normal, 0 stub, 0 nssa`
	
	`Area (0.0.0.0) (Active)`
	`Interfaces in this area: 2 Active Interfaces: 2`
	`Passive Interfaces: 0`
	`SPF Calculation has run 6 times`
	`This area is Normal area`
	`Number of LSAs: 3, checksum sum 161346`	

Related Commands	N/A	
Note		

show ip ospf border-routers

show ip ospf border-routers [vrf <vrf-name>]

Displays routing table entries to an Area Border Routers.

Syntax Description	vrf	OSPF routing table entries to an Area Border Routers on specific VRF.
Default	VRF: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF parameter and updated Example
Role	admin	

Example	`switch (config)# show ip ospf border-routers vrf myvrf`
	
	`OSPF Process ID 2, vrf myvrf Internal Routing Table`
	`Codes: i - Intra-area route, I - Inter-area route`
	`i  1.1.1.1 [0]   ABR   Area: 0.0.0.0,  Next Hop: 21.21.21.1`	

Related Commands	N/A	
Note		

show ip ospf database

show ip ospf database [summary] [<process-id> <area-id> [<link-state-id>]] [adv-router <ip-address> | self-originated] [vrf <vrf-name>]

Displays the OSPF database.

Syntax Description	adv-router <ip-address>	Filters per advertise router
	area-id	Filters the command per OSPF Area ID. Range is 0-4294967295.
	link-state-id	The link state ID
	self-originated	Self Originate
	summary	Summarizes the output of the OSPF database.
	process-id	Displays OSPF database on specific instance ID
	vrf	Displays OSPF database on specific VRF
Default	Process ID: 1 VRF: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Role	admin	
Example	`switch (config)# show ip ospf database 2 vrf myvrf`	

```
OSPF Router with ID (2.2.2.2) (Process ID 2 VRF myvrf)

                Router Link States (Area 0.0.0.0)
                ----------------------------------------
Link ID         ADV Router      Age             Seq             Checksum        LinkCount

2.2.2.2         2.2.2.2         1150            0x80000006      0xbd2a          3

1.1.1.1         1.1.1.1         1152            0x80000006      0xf7f5          3

                Network Link States (Area 0.0.0.0)
                ----------------------------------------
Link ID         ADV Router      Age             Seq             Checksum

21.21.21.2      2.2.2.2         1150            0x80000003      0xbb26
```

Related Commands	N/A
Note	

show ip ospf interface

show ip ospf interface [<process-id>] [vlan <vlan-id> | Ethernet <slot/port | port-channel <number>] [brief]

Displays the OSPF related interface configuration.

Syntax Description	brief	Gives a brief summary of the output
	process-id	Displays OSPF interface configuration on specific instance ID
	vlan <vlan-id>	Displays OSPF interface configuration and status per VLAN interface
	vrf	Displays OSPF interface configuration on specific VRF
Default	Process ID: 1 VRF: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
	3.6.4070	Added Ethernet variable
Role	admin	

Example

```
switch (config) # show ip ospf interface 2 vrf myvrf

Interface Vlan is 21 Enabled, line protocol is Up
IP address 21.21.21.2,  Mask 255.255.255.0 [primary]
IP address 30.30.30.30, Mask 255.255.255.0
Process ID 2 VRF myvrf, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DESIGNATED ROUTER, Network Type BROADCAST, Cost 10
Transmit delay 1 sec, Router Priority 1
DR is 2.2.2.2
Backup Designated Router is 1.1.1.1
Timer intervals (secs): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0

switch (config) # show ip ospf interface 2 vrf myvrf brief

OSPF Process ID 2 VRF myvrf
Total number of interface: 2
Interface Id    Area          Cost          State          Neighbors      Status
Vlan21          0.0.0.0       10            Enabled        1              Up
Ethernet1/22    0.0.0.0       1             Enabled        1              Up
```

Related Commands	N/A
Note	

show ip ospf neighbors

show ip ospf [vrf <vrf-name>] neighbors [vlan <vlan-id> | interface <name>] [<neighbor ip address>]

Displays the OSPF related interface neighbor configuration.

Syntax Description	vlan-id	Displays OSPF interface configuration and status per VLAN interface
	neighbor ip address	Filers the output per a specific OSPF neighbor
	vrf	Displays OSPF interface neighbor configuration on specific VRF
Default	VRF: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF parameter and updated Example
	3.6.4070	Added support for BFD
Role	admin	

Example

```
switch (config) # show ip ospf neighbors vrf myvrf

Neighbor 1.1.1.1, interface address 21.21.21.1
In the area 0.0.0.0  via Interface Vlan 21
Neighbor priority is 1, State is FULL
DR is  2.2.2.2
Backup Designated Router is 1.1.1.1
Options 2
Dead timer due in 36

Neighbor 1.1.1.1, interface address 22.22.22.1
In the area 0.0.0.0  via  1/22
Neighbor priority is 1, State is FULL
No designated router on this network
No backup designated router on this network
Options 2
Dead timer due in 36
switch (config) # show ip ospf neighbors  1/22 vrf myvrf

Neighbor 1.1.1.1, interface address 22.22.22.1
In the area 0.0.0.0  via  1/22
Neighbor priority is 1, State is FULL
No designated router on this network
No backup designated router on this network
Options 2
Dead timer due in 29
```

Related Commands	N/A
Note	BFD session state is displayed as: established, failed or not established. When BFD is not defined in the command, it is not displayed in the output.

show ip ospf request-list

show ip ospf request-list <neighbor-id> {vlan <vlan-id> | ethernet <slot/port> | port-channel <id>} [vrf <vrf-name>]

Displays the OSPF list of all link-state advertisements (LSAs) requested by a router.

Syntax Description	neighbor-id	Filers the output per a specific OSPF neighbor.
	vlan-id	Filers the output per a specific VLAN ID.
	vrf <vrf-name>	Displays OSPF request-list on specific VRF
Default	vrf: active VRF routing-context	
Configuration Mode	Any command mode	
History	3.6.3506 3.3.3500	
Role	admin	

Example	`switch (config) # show ip ospf request-list 4.4.4.4 vlan 7`
	
	`OSPF Router with ID (7.7.7.1) (Process ID 1)`
	`Neighbor 4.4.4.4, Interface vlan 7, Address 7.7.7.2`
	`42 LSAs on request-list`
	
	`Type LS-ID ADV-RTR Seq No Age Checksum`
	`1 10.10.10.23 10.10.10.23 0x8000012f 37 0xa7b9`
	`1 10.10.10.24 10.10.10.24 0x8000012f 38 0xbd61`	

Related Commands	N/A	
Note		

show ip ospf retransmission-list

show ip ospf retransmission-list <neighbor-id> {vlan <vlan-id> | ethernet <slot/port> | port-channel <id>} [vrf <vrf-name>]

Displays the OSPF list of all link-state advertisements (LSAs) waiting to be resent to neighbors.

Syntax Description	neighbor-id	Filers the output per a specific OSPF neighbor.
	vrf <vrf-name>	Displays OSPF retransmission-list on specific VRF.
	vlan-id	Filers the output per a specific VLAN ID.
Default	vrf: Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.6.3506 3.3.3500	
Role	admin	

Example	`switch (config) # show ip ospf retransmission-list 4.4.4.4 vlan 6`
	
	`OSPF Router with ID (7.7.7.1) (Process ID 1)`
	`Neighbor 4.4.4.4, Interface vlan 6, Address 6.6.6.2`
	`Link state retransmission due in 3780 msec, Queue length 207`
	
	`Type LS-ID ADV-RTR Seq No Age Checksum`
	`3 22.22.22.22 7.7.7.1 0x80000045 0 0xaaf4`
	`3 192.168.23.2 7.7.7.1 0x80000001 353 0x6752`	

Related Commands	N/A	
Note		

show ip ospf summary-address

show ip ospf summary-address [vrf <vrf-name>]

Displays a list of all summary address redistribution information configured on the OSPF.

Syntax Description	vrf <vrf-name>	Display summary address and area range information on specific VRF.
Default	vrf : Active VRF routing-context	
Configuration Mode	Any command mode	
History	3.6.3506 3.3.3500	
Role	admin	
Example		

```
switch (config)# show ip ospf summary-address

OSPF Process ID 1 VRF default
Network         Mask            Area            Advertise       LSA type        Metric          Tag
------          -----           -----           -------         ---------       -------         -----
66.66.66.0      255.255.255.0   0.0.0.1         Advertise       Type 3          Auto            N/A
66.66.66.0      255.255.255.0   0.0.0.1         Advertise       Type 7          Auto            N/A
55.55.55.0      255.255.255.0   0.0.0.5         Advertise       Type 3          Auto            N/A
33.33.0.0       255.255.0.0     N/A             Advertise       Type 5          Auto            N/A
44.44.0.0       255.255.0.0     N/A             Advertise       Type 5          Auto            N/A
arc-switch111 [standalone: master] (config) #
```

Related Commands	N/A
Note	

6.4 BGP

Border Gateway Protocol (BGP) is an exterior gateway protocol which is designed to transfer routing information between routers. It maintains and propagates a table of routes which designates network reachability among autonomous systems (ASs).

BGP neighbors, or peers, are routers configured manually to converse using the BGP protocol on top of a TCP session on port 179. A BGP speaker periodically sends keep-alive messages to maintain the connection. Network reachability includes such information as forwarding destinations (IPv4 or IPv6) together with a list of ASs that this information traverses and other attributes, so it becomes possible to construct a graph of AS connectivity without routing loops. BGP makes possible to apply policy rules to enforce connectivity graph.

BGP routers communicate through TCP connection on port 179. Connection between BGP neighbors is configured manually or can be established dynamically by configuring dynamic listen groups. When BGP runs between two peers in the same AS, it is referred to as Internal BGP (iBGP, or Interior Border Gateway Protocol). When it runs between separate ASs, it is called External BGP (eBGP, or Exterior Border Gateway Protocol). Both sides can initiate a connection, after the initial connectivity is created, BGP state machine drives both sides to enter into ESTABLISHED state where they can exchange UPDATE messages with reachability information.

6.4.1 State Machine

In order to make decisions in its operations with peers, a BGP peer uses a simple finite state machine (FSM) that consists of six states: Idle; Connect; Active; OpenSent; OpenConfirm; and Established. For each peer-to-peer session, a BGP implementation maintains a state variable that tracks which of these six states the session is in. The BGP protocol defines the messages that each peer should exchange in order to change the session from one state to another.

The first state is the “Idle” state. In “Idle” state, BGP initializes all resources, refuses all inbound BGP connection attempts and initiates a TCP connection to the peer. The second state is “Connect”. In the “Connect” state, the router awaits the TCP connection to complete and transitions to the “OpenSent” state if successful. If unsuccessful, it initializes the ConnectRetry timer and transitions to the “Active” state upon expiration. In the “Active” state, the router resets the ConnectRetry timer to zero and returns to the “Connect” state. In the “OpenSent” state, the router sends an Open message and waits for one in return in order to transition to the “OpenConfirm” state. KeepAlive messages are exchanged and, upon successful receipt, the router is placed into the “Established” state. In the “Established” state, the router can send/receive: KeepAlive; Update; and Notification messages to/from its peer.

6.4.2 Default Address Family

Default Address Family defines which address family is activated when peer or peer-group becomes active.

When the default address family configuration is modified - it will cause a renegotiation of capabilities for all neighbors that do not have explicit configuration of active address families.

The default address family in BGP is IPv4.

6.4.3 Default Route Originate

Default Route Originate initial value is set to “false”.

6.4.4 Peer Groups and Update Groups

Any BGP peer can be defined as part of a peer group and it will inherit peer group configuration or have its own configuration.

A system will automatically generate an update group from peer groups members.

Peer that has a different outbound policy from peer-group will not become a part of update group.

6.4.5 Configuring BGP

Figure 40: Basic BGP Configuration

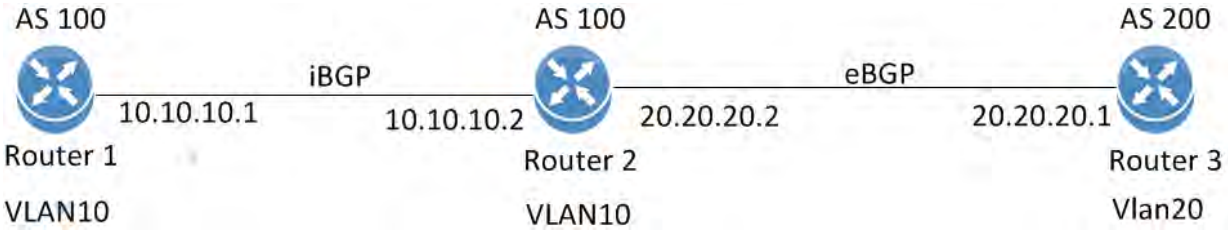

Follow these steps for basic BGP configuration on two switches (Router 1 and Router 2):

Preconditions:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:

```
switch (config)# vlan 10
```

The same VLAN must be configured on both switches.

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config ethernet 1/1)# switchport access vlan 10
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

Step 5. Apply IP address to the VLAN interface on Router 1. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.1 /24
```

Step 6. Apply IP address to the VLAN interface on Router 2. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /24
```

Step 7. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

Configure BGP:

Step 1. Enable BGP. Run:

```
switch (config)# protocol bgp
```

Step 2. Configure an AS number that identifies the BGP router. Run:

```
switch (config)# router bgp 100
```

To run iBGP, the AS number of all remote neighbors should be identical to the local AS number of the configured router.

Step 3. Configure BGP Router 1 neighbor. Run:.

```
switch (config router bgp 100)# neighbor 10.10.10.2 remote-as 100
```

Step 4. Configure BGP Router 2 neighbor. Run:.

```
switch (config router bgp 100)# neighbor 10.10.10.1 remote-as 100
```

6.4.6 Verifying BGP

Step 1. Check the general status of BGP. Run:

```
switch (config)# show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 100
BGP table version is 100, main routing table version 100
0 network entries using 0 bytes of memory
0 path entries using 0 bytes of memory
0 BGP AS-PATH entries using 0 bytes of memory
0 BGP community entries using 0 bytes of memory
0 BGP extended community entries using 0 bytes of memory
Neighbor        V       AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down     State/PfxRcd
10.10.10.2      0      100     100      76       3    0    0 00:0:10:19 ESTABLISHED
switch (config)#
BGP summary information for VRF default, address family IPv4
```

- Verify that the state of each BGP neighbor reached to ESTABLISHED state.
- In case the neighbor is disabled (shutdown). The state of the neighbor will be IDLE.
- BGP incoming and outgoing messages should be incremented.
- The AS number of each neighbor is the correct one.

Step 2. Check the status of the neighbors. Run:

```
switch (config)# show ip bgp neighbors
BGP neighbor is 10.10.10.2, remote AS 100, external link
   BGP version 0, remote router ID 0.0.0.0
   BGP State = ESTABLISHED
   Last read 0:00:00:00, last write 0:00:00:00, hold time is 180, keepalive
interval is 60 seconds
   Configured hold time is 180, keepalive interval is 60 seconds
   Minimum holdtime from neighbor is 0 seconds
switch (config)#
```

You should be able to see running BGP counters and ESTABLISHED state per active neighbor.

6.4.7 Ethernet Virtual Private Network

Ethernet Virtual Private Network (EVPN) technology provides L2 and L3 VPN services by advertising Ethernet MAC addresses and IP routes over BGP address family. This technology supports multiple forwarding planes including VXLAN.

BGP L2-EVPN address family distributes EVPN “routes” between EVPN enabled nodes where some of them are Virtual Tunnel Endpoints (VTEPs) with VXLAN functionality and some of them are transit nodes that perform BGP reflection functionality.

The following route types are defined by RFC 7432:

- MAC/IP advertisement route (route type 2) – advertises MAC and IP addresses of end-systems and their mapping to broadcast domains (VXLAN VNIs and EVPN EVIs). It is used for unicast forwarding, ARP suppression, and advertising default gateway in the EVPN network.
- Inclusive multicast Ethernet tag route (route type 3) – advertises EVPN bridge domain (EVI) and originating router IP address. The EVPN network uses those addresses to instantiate forwarding plane for BUM (Broadcast, unknown Unicast, unknown Multicast) traffic.
- IP prefix route (type 5) – advertises IP prefix, IP gateway, IP address, and HW encapsulation (VNI in the case of VXLAN). This route is used to establish IP prefix LPM routing in the EVPN nodes.

Other route types (type 1 and 4) are used in multi-homing environments only.

RFC 7432 defines BGP attributes that should be used together with L2-EVPN address family routes:

- PMSI tunnel attributes – used for inclusive multicast Ethernet tag route to define multicast type (head end replication) and data path (VNI)
- MAC mobility extended community – used in MAC/IP routes to inform neighbors about MAC roaming events
- Default gateway – used by MAC/IP route to establish default gateway routes
- Route targets – used by all routes to import and export BGP L2-VPN to forwarding and from plane

6.4.8 BGP Commands

Some of the commands in Section 6.4.8.2, “Config Router,” on page 1382 feature the parameters “no” and “disable”.

The parameter “no”:

- Removes the command from running-config
- If configured value can be inherited, it is inherited
- If not inherited and has default, it takes the default
- If used on a string value, that value is removed (e.g. “password”, “no password”) or inherited
- If used on a numerical value, the default value is taken
- If used on a boolean value, it either takes default or FALSE, or inherits value

The parameter “disable”:

- Creates an entry in running-config
- Prevents inheritance
- If used on a string value, that value is removed (e.g. “password”, “no password”, cannot be inherited)
- If used on a numerical value, it sets a default numeric value (e.g. “disable neighbor 1.1.1.1 weight” sets the default weight)
- If used on a boolean value, it sets value to FALSE (e.g. “disable neighbor 1.1.1.1 send-community” disables send-community, and its configuration cannot be inherited)

6.4.8.1 Config

protocol bgp

	protocol bgp **no protocol bgp** Enables BGPv4, and unhides BGP related commands. The no form of the command deletes all BGP configuration and hides BGP related commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.5006
Role	admin
Example	`switch (config)# protocol bgp` `switch (config)#`
Related Commands	ip routing
Note	

clear ip bgp

clear ip bgp [{<ip-address> | all} [soft] [in]]

Clears BGP learned routes from the BGP table and resets the connection to the neighbor.

Syntax Description	ip-address	A BGP peer IP address. Only the specified neighbor is reset.
	all	All BGP peers. All BGP neighbors are reset.
	soft	Clears BGP learned routes from the BGP table without resetting the connection to the neighbor
	in	Inbound routes are reset
Default	N/A	
Configuration Mode	config	
History	3.3.5006	First release
	3.3.5200	Updated description
	3.6.3004	Removed “out” parameter
Role	admin	
Example	`switch (config)# clear ip bgp all` `switch (config)#`	
Related Commands	N/A	
Note	This command removes BGP IPv4 learned routes from the routing table, reads all routes from designated peers, and sends routes to those peers as required.	

router bgp

router bgp <as-number>
no router bgp <as-number>

Creates and enters a BGP instance with the specified AS number.
The no form of the command deletes all router BGP instance configuration.

Syntax Description	as-number	Autonomous system number: A unique number to be used to identify the AS. The AS is a number which identifies the BGP router to other routers and tags the routing information passed along. Range: 1-(2^{32}-1).
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
	3.3.5200	Updated syntax description
Role	admin	
Example	`switch (config)# router bgp 100` `switch (config router bgp 100)#`	
Related Commands	ip routing	
Note		

6.4.8.2 Config Router

shutdown

	shutdown **no shutdown** Gracefully disables BGP protocol without removing existing configuration. The no form of the command enables BGP.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config router bgp
History	3.3.5006
Role	admin
Example	`switch (config router bgp 100)# no shutdown`
Related Commands	
Note	

address-family

address-family {ipv4-unicast | ipv6-unicast | l2vpn-evpn}

Enables selected address family configuration mode.

Syntax Description	ipv4-unicast	Enables IPv4 address family configuration mode
	ipv6-unicast	Enables IPv6 address family configuration mode
	l2vpn-evpn	Enables EVPN address family configuration mode
Default	IPv4	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.8100	Added “l2vpn-evpn” parameter
Role	admin	
Example	`switch (config router bgp 65001) # address-family l2vpn-evpn` `switch (config router bgp 65001 address-family l2vpn-evpn) #`	
Related Commands		
Note		

aggregate-address

aggregate-address <ip_prefix_length> [summary-only] [as-set] [attribute-map]
no aggregate-address <ip_prefix_length> [summary-only] [as-set] [attribute-map]

Creates an aggregate route in the BGP database.
The no form of the command disables ECMP across AS paths.

Syntax Description	ip_prefix_length	Destination to aggregate
	summary-only	Contributor routes are not advertised
	as-set	Includes AS_PATH information from contributor routes as AS_SET attributes
	attribute-map	Assigns attribute values in set commands of the map's permit clauses. Deny clauses and match commands in permit clauses are ignored.
Default	Disabled	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added support for IPv4 and IPv6
Role	admin	
Example	`switch (config router bgp 4) # aggregate-address 3.5.3.7 /32`	
Related Commands		
Note	• Aggregate routes combine the characteristics of multiple routes into a single route that the switch advertises • Aggregation can reduce the amount of information that a BGP speaker is required to store and transmit when advertising routes to other BGP speakers • Aggregate routes are advertised only after they are redistributed	

bestpath as-path multipath-relax

bestpath as-path multipath-relax [force]
no bestpath as-path multipath-relax [force]

Enables ECMP across AS paths.
The no form of the command disables ECMP across AS paths.

Syntax Description	force	Applies configuration while BGP is admin-up
Default	Disabled	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description and notes
	3.6.3004	Added “force” parameter
Role	admin	
Example	`switch (config router bgp 100)# bestpath as-path multipath-relax`	
Related Commands	maximum-paths	
Note	• With this option disabled, only routes with exactly the same AS path as the best route to a destination are considered for ECMP • With this option enabled, all routes with similar length AS path as the best route are considered for ECMP	

bgp default

bgp default {ipv4-unicast | ipv6-unicast}
no bgp default {ipv4-unicast | ipv6-unicast}

Enables setting address families as default for peer or peer-group activation.
The no form of the command disables setting address families as default for peer or peer-group activation.

Syntax Description	ipv4-unicast	IPv4 unicast address family, enabled by default.
	ipv6-unicast	IPv6 unicast address family, disabled by default.
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.4110	Added support for IPv6
Role	admin	
Example	`switch (config router bgp 100)# bgp default ipv4-unicast`	
Related Commands		
Note	This command can be used multiple times and each address family can be configured separately.	

bgp fast-external-fallover

	bgp fast-external-fallover **no bgp fast-external-fallover** Terminates eBGP sessions of any directly adjacent peer without waiting for the hold-down timer to expire if the link used to reach the peer goes down. The no form of the command waits for hold-down timer to expire before terminating eBGP sessions.
Syntax Description	N/A
Default	no bgp fast-external-fallover
Configuration Mode	config router bgp
History	3.4.0000
Role	admin
Example	`switch (config router bgp 100)# bgp fast-external-fallover`
Related Commands	maximum-paths
Note	Although this feature improves BGP conversion time, it may cause instability in your BGP table due to a flapping interface.

bgp listen limit

bgp listen limit <maximum>
no bgp listen limit

Limits the number of dynamic BGP peers allowed on the switch.
The no form of the command resets to the default value.

Syntax Description	maximum	The maximum number of dynamic BGP peers to be allowed on the switch. Range: 1-128.
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
Role	admin	
Example	`switch (config router bgp 100)# bgp listen limit 101`	
Related Commands		
Note		

bgp listen range peer-group

	bgp listen range <ip_prefix> peer-group <peer-group-name> remote-as <as-number> **no bgp listen range <ip_prefix> <length>** Identifies a range of IP addresses from which the switch will accept incoming dynamic BGP peering requests. After applying the no form of the command, the switch will no longer accept dynamic peering requests on the range.	
Syntax Description	ip-address	IP address
	length	Mask length (e.g. /24 or 255.255.255.254)
	peer-group-name	Peer group name
	remote-as <as-number>	Remote peer's number
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Added note
Role	admin	
Example	`switch (config router bgp 100)# bgp listen range 10.10.10.10 /24 peer-group my-group remote-as 13`	
Related Commands		
Note	• To create a static peer group, use the command `neighbor peer-group` • Neighbors in a dynamic peer group are configured as a group and cannot be configured individually • The no form of the command may take up to a few seconds to take effect if there are many dynamic peers and/or a lot of routes. While the clean-up process is running, creation of a new listen range that overlaps the deleted one will fail. • If dynamic peer range is defined with an overlap to another defined range, the longest remote address prefix take affect	

cluster-id

cluster-id <ip-address> [force]
no cluster-id <ip-address> [force]

Configures the cluster ID in a cluster with multiple route reflectors.
The no form of the command resets the cluster ID for route reflector.

Syntax Description	ip-address	The route reflector cluster ID • 0.0.0.1 to 255.255.255.255 Valid cluster ID number • 0.0.0.0 removes the cluster-ID from the switch (similar to “no cluster-id”)
	force	Applies configuration while BGP is admin-up
Default	Cluster ID is the same as Router ID	
Configuration Mode	config router bgp	
History	3.2.1000	
	3.4.0000	Updated syntax description
	3.6.3004	Added “force” parameter
Role	admin	
Example	`switch (config router bgp 100)# cluster-id 10.10.10.10`	
Related Commands	N/A	
Note		

client-to-client reflection

	client-to-client reflection **no client-to-client reflection** The switch will be configured as a route reflector. The no form of the command stops the switch from being a route reflector
Syntax Description	N/A
Default	client-to-client reflection is enabled
Configuration Mode	config router bgp
History	3.2.1000
Role	admin
Example	`switch (config router bgp 100)# client-to-client reflection`
Related Commands	N/A
Note	

distance

distance <external> <internal> <local>
no distance

Sets the administrative distance of the routes learned through BGP.
The no form of the command resets the administrative distance its default.

Syntax Description	external	Administrative distance for external BGP routes. Range: 1-255.
	internal	Administrative distance for internal BGP routes. Range: 1-255.
	local	Administrative distance for local BGP routes. Range: 1-255.
Default	external: 20 internal: 200 local: 200	
Configuration Mode	config router bgp	
History	3.3.5006	
Role	admin	
Example	`switch (config router bgp 100)# distance 10 20 30`	
Related Commands	N/A	
Note	• Routers use administrative distances to decide on a route when two protocols provide routing information to the same destination. • Lower distance values correspond to higher reliability. • Routes are external when learned from an external autonomous system. • Routes are internal when learned from a peer in the local autonomous system. • Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks being redistributed from another process. • BGP routing tables do not include routes with a distance of 255.	

graceful-restart stalepath-time

graceful-restart stalepath-time <interval>
no graceful-restart stalepath-time

Configures the maximum time that stale routes from a restarting BGP neighbor are retained after a BGP session is reestablished with that peer.
The no form of the command resets to the default value.

Syntax Description	interval	Time in seconds. Range: 1-3600.
Default	300 seconds	
Configuration Mode	config router bgp	
History	3.4.0000	
Role	admin	
Example	`switch (config router bgp 100)# graceful-restart stalepath-time 350`	
Related Commands	N/A	
Note		

maximum-paths

maximum-paths [ibgp] <maximum-path>

Configures the maximum number of parallel eBGP/iBGP routes that the switch installs in the routing table.

Syntax Description	ibgp	Sets the configuration on the internal BGP.
	maximum-path	The number of routes to install to the routing table. Range: 1-32
Default	1	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description and notes
	3.6.4070	Updated maximum-path range
Role	admin	
Example	`switch (config router bgp 100)# maximum-paths ibgp 10` `switch (config router bgp 100)#`	
Related Commands	N/A	
Note	• This command provides an ECMP parameter that controls the number of equal-cost paths that the switch installs in the routing table for each destination. • The action is effective after BGP restart. • If the parameter “ibgp” is not used, the setting is applied on routes learned from peers from other ASs; if “ibgp” is used, the setting is applied to routes learned from peers of the same AS.	

neighbor activate

neighbor <ip-address | peer-group> activate
no neighbor <ip-address | peer-group> activate
disable neighbor <ip-address | peer-group> activate

Sends advertisement for given address-family to neighbor.
The no form of the command removes the command from running-config and enables inheritance.
The disable form of the command sets boolean value to false, and disables inheritance.

Syntax Description	ip-address	Neighbor IP address
	peer-group	Peer group name
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.6.4070	
	3.6.4110	Added “disable” option to the command
	3.6.8100	Added “config router bgp address-family” configuration mode
Role	admin	
Example	`switch (config router bgp 100)# no neighbor 10.10.10.1 activate` `switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.3.2 activate`	
Related Commands		

Note	There are 4 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 activate`
	• At the end of the command
	`switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 activate disable`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 activate`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 activate`	

neighbor advertisement-interval

neighbor {<ip-address> | <peer-group-name>} advertisement-interval <delay>
no neighbor {<ip-address> | <peer-group-name>} advertisement-interval

Sets the minimum route advertisement interval (MRAI) between the sending of BGP routing updates.
The no form of the command disables this function.

Syntax Description	ipv4_addr, ipv6_addr	A BGP peer IP address
	peer-group-name	Peer group name
	delay	Time (in seconds) is specified by an integer Range: 0-600 where "0" disables this function and prevents the system from inheriting this parameter's group configuration
Default	30 seconds	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "delay" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 advertisement-interval 100`	
Related Commands		
Note	When configuring an advertisement interval to a BGP session, this interval is implemented per prefix route of that session. For example: If a session is configured with advertisement interval of 100 seconds, when it first learns a new route it automatically sends an update on this route. If it learns another route in the same prefix as the initial route, it waits for 100 seconds. But if it learns another route in a different prefix it immediately advertises that route and does not wait another 100 seconds.	

neighbor allowas-in

neighbor {<ip-address > | <peer-group-name>} allowas-in [number]
no neighbor {<ip-address > | <peer-group-name>} allowas-in

Configures the switch to permit the advertisement of prefixes containing duplicate autonomous switch numbers (ASNs).
The no form of the command disables this function.

Syntax Description	ip-address	A BGP peer IP address
	peer-group-name	Peer group name
	number	Number of switch's (ASN) allowed in path Range: 0-10 where "0" disables this function and prevents the system from inheriting this parameter's group configuration
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "number" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 allowas-in 2`	
Related Commands	ip routing router bgp <as-number>	
Note	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers.	

neighbor default-originate

[no | disable] neighbor <ip-address | peer_group> default-originate [route_map_name]

Enables advertisement of the default route to a specified neighbor or peer group. The no form of the command disables advertisement of the default route.

Syntax Description	ip-address	Neighbor IPv4 address
	peer_group	Peer group's name
	route_map_name	route map name that modifies default route attributes
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.4110	Added "disable" option to the command
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.1 default-originate default-attr`	
Related Commands		
Note		

neighbor description

neighbor {<ip-address> | <peer-group-name>} description <string>
no neighbor {<ip-address> | <peer-group-name>} description

Associates descriptive text with the specified peer or peer group.
The no form of the command removes the description from the peer.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	string	Free string, up to 80 characters in length
Default	No description	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated example
	3.6.4070	Added support for IPv6 and IPv4
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 description The next door neighbor`	
Related Commands	N/A	
Note	The peer description only appears in the show commands.	

neighbor ebgp-multihop

neighbor {<ip-address > | <peer-group-name>} ebgp-multihop [<ttl>]
no neighbor {<ip-address > | <peer-group-name>} ebgp-multihop

Enables BGP to connect to external peers that are not directly connected to the switch. The no form of the command disables connecting to external peers.

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	ttl	Time-to-live Range: 1-255 hops where "1" disables connecting to external peers and prevents the system from inheriting this parameter's group configuration
Default	ttl: 1	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated default
	3.6.3004	Updated description of "ttl" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 ebgp-multihop 5`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Note	The command does not establish the multi-hop if the only route to the peer is the default route (0.0.0.0).	

neighbor export-localpref

neighbor {<ip-address> | <peer-group-name>} export-localpref <value>
no neighbor {<ip-address> | <peer-group-name>} export-localpref

Configures the local preference value sent to the specified peer or peer group. The no form of the command resets the local preference to its default value.

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	value	Preference value Range: 0-2147483647 where "100" configures the default, and prevents the system from inheriting this parameter's group configuration
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "value" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 export-localpref 100`	
Related Commands		
Note		

neighbor fall-over bfd

[no] neighbor {<ip-address> | <ip-address> | <peer-group-name>} fall-over bfd

Disables BFD as a mechanism to detect failure.
The no form of the command enables BFD neighbor.

Syntax Description	peer-group-name	Peer group name
	ip-address	IP address of the neighbor
Default	Enabled	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 bfd`	
Related Commands		
Note	The command “no neighbor <ip_address> fall-over bfd” affects traffic, BGP will restore the connection based on Hello protocol.	

neighbor import-localpref

[no] neighbor {<ip-address> | <peer-group-name>} import-localpref <value>

Configures the local preference value assigned to routes received from the specified peer or peer group.
The no form of the command resets the local preference to its default value.

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	value	Preference value Range: 0-2147483647 where “100” configures the default, and prevents the system from inheriting this parameter’s group configuration
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “value” parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 import-localpref 100`	
Related Commands		
Note		

neighbor local-as

neighbor {<ip-address> | <peer-group-name>} local-as <asn-id> [no-prepend | replace-as]
no neighbor {<ip-address> | <peer-group-name>} local-as

Enables the modification of the AS path attribute for routes received from an eBGP neighbor.
The no form of the command disables AS path modification for the specified peer or peer group.

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	asn-id	AS number that is sent instead of the actual AS of the switch. Range: 0-4294967295
	no-prepend	local-as number is not pre-pended to the routes received from external neighbors
	replace-as	Prepends only the local autonomous system number (as configured with the IP address argument) to the AS path attribute
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "as-id" parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.4110	Updated command.
Role	admin	
Example	`switch (config router bgp 4) # neighbor 100.100.100.100 local-as 123`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Note	• This function allows the switch to appear as a member of a different autonomous system (AS) to external peers. • To disable peering with the neighbor run the command `clear ip bgp`	

neighbor local-v6-addr

neighbor {<ip-address > | <peer-group-name>} local-v6-addr <ipv6_local>
no neighbor {<ip-address > | <peer-group-name>} local-v6-addr

Specifies the switche's next-hop value sent using IPv6 NLRI in IPv4 transport session.
The no form of the command removes next-hop value.

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	ipv6_local	IPv6 next hop address
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 4) # neighbor 10.10.10.1 local-v6-addr 2001::2`	
Related Commands		
Note		

neighbor maximum-prefix

neighbor {<ip-address> | <peer-group-name>} maximum-prefix <maximum> [warning-only]
no neighbor {<ip-address> | <peer-group-name>} maximum-prefix

Configures the number of BGP routes the switch accepts from a specified neighbor and defines an action when the limit is exceeded.
The no form of the command removes the limitation

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	maximum	Number of BGP routes the switch accepts from a specified neighbor Range: 1-2147483647 where "12000" configures the default, and prevents the system from inheriting this parameter's group configuration
	warning-only	Only generates a warning rather than disconnecting the neighbor
Default	12000	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "maximum" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 maximum-prefix 12000 warning-only`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Note		

neighbor next-hop-peer

neighbor {<ip-address> | <peer-group-name>} next-hop-peer [disable]
no neighbor {<ip-address> | <peer-group-name>} next-hop-peer

Configures the switch to list the peer address as the next hop in routes that it receives from the specified peer BGP-speaking neighbor or members of the specified peer group.
The no form of the command disables this function.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Disables this function and prevents the system from inheriting this parameter's group configuration
Default	no next-hop-peer	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.3004	Added "disable" parameter
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-peer`	
Related Commands		
Note	This command overrides the next hop for all routes received from this neighbor or peer group	

neighbor next-hop-self

neighbor {<ip-address> | <peer-group-name>} next-hop-self [disable]
no neighbor {<ip-address> | <peer-group-name>} next-hop-self

Configures the IP address of the router as the next hop address in routes advertises to the specific neighbor.
The no form of the command resets this parameter to its default.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Disables this function and prevents the system from inheriting this parameter's group configuration
Default	no next-hop-self	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.4070	Added support for IPv6
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-self`	
Related Commands	neighbor <ip-address> remote-as <as-number>	
Note	• This function is used in networks where BGP neighbors do not directly access all other neighbors on the same subnet. • In the default state, the next hop is generated based on the IP address and the present next hop in the route information.	

[neighbor] next-hop-unchanged

[neighbor <ip-address | peer group>] next-hop-unchanged
[no | disable] [neighbor <ip-address | peer group>] next-hop-unchanged

Enables preserving BGP next-hop when forwarding routes to this eBGP peer or all eBGP peers in this address family.
The no form of the command removes configuration and enables inheritance of AFI SAFI next-hop-unchanged configuration from a peer group if this neighbor is member in one.
The disable form of the command disables preserving BGP next-hop when forwarding routes to this eBGP peer or all eBGP peers in this address family.

Syntax Description	ip-address	Neighbor IP address
	peer_group	Peer group name
Default	By default, the next-hop of a route is preserved when advertising the route to an iBGP peer, but is updated when advertising the route to an eBGP peer. Setting this to “true” overrides this behavior and preserves the next-hop when routes are advertised to this eBGP peer.	
Configuration Mode	config router bgp address-family	
History	3.6.8100	
Role	admin	
Example	`switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.5.2 next-hop-unchanged` `switch (config router bgp 65001 address-family l2vpn-evpn) # next-hop-unchanged`	
Related Commands	address-family l2vpn-evpn	

Note	There are 4 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged`
	• At the end of the command
	`switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged disable`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 next-hop-unchanged`	

neighbor password

neighbor {<ip-address> | <peer-group-name>} password [<encryption>] <string>
no neighbor {<ip-address> | <peer-group-name>} password

Enables authentication on a TCP connection with a BGP peer.
The no form of the command resets the value to its default.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	encryption	Possible values: • no parameter – clear text • 0 – clear text • 7 – obfuscated
	string	Up to 8 bytes in length
Default	no neighbor password	
Configuration Mode	config router bgp	
History	3.4.0000	
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 password 7 admin123`	
Related Commands		
Note	• Peers must use the same password to ensure communication. • neighbor <ip-address> password 7 <password>' can only accept data that was created using 'show config'. • 'show config' will never show the clear-test password, it will always be obfuscated (and thus displayed using the 'password 7' syntax). • Router BGP neighbor password cannot be set when enabling secure mode • Router BGP peer-group password cannot be set when enabling with secure mode	

neighbor no-password

neighbor {<ip-address> | <peer-group-name>} no-password

Disables authentication for peer without inheritance.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.3004	
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 no-password`	
Related Commands	neighbor password	
Note		

neighbor peer-group

1. neighbor {<ip-address >} peer-group <peer-group-name>
2. neighbor {<peer-group-name>} peer-group
3. no neighbor {<ip-address >} peer-group <peer-group-name>
4. no neighbor {<peer-group-name>} peer-group

1. Assigns BGP neighbors to an existing peer group
2. Creates a peer-group
3. Unassigns a BGP neighbor from a peer-group
4. Deletes the peer-group

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Added notes
	3.6.4070	Added support for IPv6 and IPv4
Role	admin	
Example	`switch (config router bgp 100)# neighbor groupA peer-group` `switch (config router bgp 100)# neighbor 1.2.3.4 peer-group groupA`	
Related Commands		
Note	• Once a peer group is created, the group name can be used as a parameter in neighbor configuration commands, and the configuration will be applied to all members of the group. • Settings applied to an individual neighbor in the peer group override group settings. • A neighbor can only belong to one peer group, so issuing this command for a neighbor that is already a member of another group removes it from that group. • When a neighbor is removed from a peer group, the neighbor retains the configuration inherited from the peer group. • Router BGP peer-group password cannot be set when enabling with secure mode • A BGP group must be used by either a single listen range, or by a set of neighbors sharing the same type (iBGP or eBGP) • A group must already exist before a node is configured to use it • Any configuration change on a group affects each of the peers inheriting this specific parameter from the group only after undergoing admin state toggle	

neighbor remote-as

neighbor {<ip-address>} remote-as <as-number>
no neighbor {<ip-address>} remote-as <as-number>

Configures a neighbor.
The no form of the command removes the neighbor, dropping the connection and all routes if already connected.

Syntax Description	ipv4_addr, ipv6_addr	IP address of the neighbor
	peer-group-name	Peer group name
	as-number	The BGP peer as-number. Range: 1-65535.
Default	N/A	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description and note
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 remote-as 200` `switch (config router bgp 100)#`	
Related Commands	ip routing router bgp <as-number>	
Note	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers.	

neighbor remove-private-as

neighbor {<ip-address> | <peer-group-name>} remove-private-as [disable]
no neighbor {<ip-address> | <peer-group-name>} remove-private-as

Removes private autonomous system numbers from outbound routing updates for external BGP (eBGP) neighbors.
The no form of the command preserves private AS numbers for the specified peer.

Syntax Description	ipv4_addr, ipv6_addr	A BGP peer IP address
	peer-group-name	Peer group name
	disable	Preserves private AS numbers for the specified peer and prevents the system from inheriting this parameter's group configuration
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added support for IPv6 and IPv4
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 remove-private-as` `switch (config router bgp 100)#`	
Related Commands	ip routing router bgp <as-number>	
Note	• This can only be used with external BGP (eBGP) peers. • If the update has only private AS numbers in the AS path, BGP removes these numbers. • If the AS path includes both private and public AS numbers, BGP does not remove the private AS numbers. This situation is considered a configuration error. • If the AS path contains the AS number of the eBGP neighbor, BGP does not remove the private AS number. • If the AS path contains confederations, BGP removes the private AS numbers only if they come after the confederation portion of the AS path.	

neighbor route-map

neighbor <ip-address | peer-group-name> route-map <route-map-name> [in | out]
[no | disable] neighbor <ip-address | peer-group-name> route-map <route-map-name> [in | out]

Configures route-map export or import to the peer either for a specific address family or for all (depending on the configuration context).
The no form of the command removes map-route configuration and enables inheritance. The Onyx inheritance priority is as follows:
a. Peer AFI-SAFI
b. Peer
c. Peer Group AFI-SAFI
d. Peer Group
The "disable" form of the command resets the route-map configuration to the default and disables inheritance.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	route-map-name	The name of the route-map
	in | out	• in – sets route import to the peer for this AFI/SAFI • out – sets route export to the peer for this AFI/SAFI If no parameter is explicitly used, both in and out are configured.
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.3.5006	
	3.3.5200	Updated notes and default
	3.4.1100	Added "out" parameter
	3.6.3004	Added note
	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added "config router bgp address-family" configuration mode
Role	admin	

Example	`switch (config router bgp 100)# neighbor 10.10.10.10 route-map MyRouteMap in`
	`switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.3.2 route-map routeMapSample in`
	`switch (config router bgp 100 address-family ipv4-unicast) # neighbor 1.1.1.1 route-map sampleRoutemap in`	

Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny | permit] [sequence-number] clear ip bgp {<ip-address> | all}

Note	There are 3 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-map`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 route-map`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 route-map`
	When inheritance is enabled (by default or when using the no form of the command), then if there is no peer AFI SAFI route-map configuration, then Onyx checks whether a route-map was at the peer level or not. If yes, then Onyx takes it. Otherwise, Onyx continues looking to the peer group AFI SAFI, and then the peer group (if a peer is member of a peer group).
	• Only one inbound route-map can be applied to a given neighbor
	• If a new route-map is applied to a neighbor, it replaces the previous route map
	• Changing a route-map only takes effect on routes received or sent after the change
	• A route-map must already exist before a node is configured to use it

neighbor no-route-map

neighbor {<ip-address> | <peer-group-name>} no-route-map

Unsets route-map for neighbor and prevents the system from inheriting this parameter's group configuration.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.3004	
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 no-route-map`	
Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny | permit] [sequence-number]	
Note		

neighbor route-reflector-client

neighbor <ip-address | peer-group> route-reflector-client
[no | disable] neighbor <ip-address | peer-group>] route-reflector-client

Configures a given peer to be a reflector client of this router for this address-family.
The no form of the command removes configuration and enables inheritance of AFI/SAFI route-reflector-client configuration from a peer group if this neighbor is member in one.
The disable form of the command removes a given peer from being a reflector client of this router for this AFI/SAFI and disables configuration inheritance.

Syntax Description	ip-address	Neighbor IP address
	peer-group	Peer group name
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.3.5006	
	3.3.5200	Updated notes and default
	3.6.3004	Added “disable” parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added “config router bgp address-family” configuration mode
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 route-reflector-client`	
Related Commands		

Note	There are 4 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client`
	• At the end of the command
	`switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client disable`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 route-reflector-client`	

neighbor send-community

neighbor <ip-address | peer group> send-community [extended]
[no | disable] neighbor <ip-address | peer group> send-community [extended]

Enables sending UPDATE messages to the peer containing BGP community attributes either for this address family or all relevant address-families.
The no form of the command removes configuration and enables inheritance of send-community attribute configuration.
The disable form of the command disables sending UPDATE messages containing BGP community attributes.

Syntax Description	ip-address	Neighbor IP address
	peer_group	Peer group name
	extended	Enables sending UPDATE messages to the peer for this address family containing extended BGP community attributes
Default	Enabled	
Configuration Mode	config router bgp config router bgp address-family	
History	3.4.0000	
	3.6.3004	Added “disable” parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added “config router bgp address-family” configuration mode
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 send-community`	
	`switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.3.2 send-community`	

Related Commands	N/A

Note	There are 4 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 send-community`
	• At the end of the command
	`switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 send-community disable`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 send-community`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 send-community`

neighbor shutdown

neighbor {<ip-address> | <peer-group-name>} shutdown [disable]
no neighbor {<ip-address> | <peer-group-name>} shutdown

Disables BGP neighbor gracefully.
The no form of the command enables BGP neighbor.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Enables BGP neighbor and prevents the system from inheriting this parameter's group configuration
Default	Enabled	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated note
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 shutdown`	
Related Commands	N/A	
Note	• Disabling a neighbor terminates all its active sessions and removes associated routing information • A group's shutdown immediately impacts every peer in this group, making them inherit this parameter	

neighbor soft-reconfiguration

neighbor {<ip-address> | <peer-group-name>} soft-reconfiguration inbound
no neighbor {<ip-address> | <peer-group-name>} soft-reconfiguration

Enables neighbor soft reconfiguration.
The no form of the command disables neighbor soft reconfiguration.

Syntax Description	peer-group-name	Peer group name
	ip-address	IP address of the neighbor
Default	Enabled	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.1 soft-reconfiguration inbound`	
Related Commands		
Note		

neighbor soft-reconfiguration inbound

neighbor <ip-address | peer group> soft-reconfiguration inbound
no neighbor <ip-address | peer group> soft-reconfiguration inbound

Enables neighbor soft reconfiguration.
The no form of the command disables neighbor soft reconfiguration.

Syntax Description	ip-address	Neighbor IPv4 address
	peer_group	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.8100	
Role	admin	
Example	`switch (config router bgp 65001) # neighbor 192.168.3.2 soft-reconfiguration inbound`	
Related Commands		
Note	This command is mandatory to show received EVPN for this neighbor.	

neighbor timers

neighbor {<ip-address> | <peer-group-name>} timers <keep-alive> <hold-time>
no neighbor {<ip-address> | <peer-group-name>} timers

Configures the keepalive and hold times for a specified peer.
The no form of the command resets the parameters to their default values.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	keep-alive	The period between the transmission of consecutive keep-alive messages • Range: 1-3600 seconds • "0" means that keepalive is not sent and the connection does not expire • Explicitly configuring the default, "60", prevents the system from inheriting this parameter's group configuration
	hold-time	The period the switch waits for a keepalive or update message before it disables peering • Range: 3-7200 seconds • "0" means that keepalive is not sent and the connection does not expire • Explicitly configuring the default, "180", prevents the system from inheriting this parameter's group configuration
Default	keep-alive: 60 seconds hold-time: 180 seconds	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description
	3.6.3004	Updated "hold-time" and "keep-alive" parameter's syntax description
	3.6.4070	Added IPv6 and IPv4 support
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 timers 65 195`	
Related Commands	neighbor <ip-address> remote-as <as-number>	
Note	Hold time must be at least 3 seconds and should be three times longer than the keep-alive setting.	

neighbor transport connection-mode passive

	neighbor {<ip-address> | <peer-group-name>} transport connection-mode passive [disable] **no neighbor {<ip-address> | <peer-group-name>} transport connection-mode passive** Sets the TCP connection for the specified BGP neighbor or peer group to passive mode. The no form of the command sets the specified BGP neighbor or peer group to active connection mode.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Sets the specified BGP neighbor or peer group to active connection mode and prevents the system from inheriting this parameter's group configuration
Default	TCP sessions initiated	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Added "disable" parameter
	3.6.4070	Added IPv6 and IPv4 support
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 transport connection-mode passive`	
Related Commands		
Note	• When the peer's transport connection mode is set to passive, it accepts TCP connections for BGP, but does not initiate them • BGP peers in active mode can both accept and initiate TCP connections for BGP	

neighbor update-source

neighbor <ip-address> update-source {ethernet <slot/port> | loopback <number> | port-channel <number> | vlan <vlan-id>}
no neighbor <ip-address> update-source

Configures the source-address for routing updates and to establish TCP connections with peers.
The no form of the command disables configured source-address for routing updates and for TCP connection establishment with a peer.

Syntax Description	ip-address	IP address of the neighbor
	ethernet <slot/port>	Ethernet interface
	loopback <number>	Loopback interface number
	vlan <vlan-id>	VLAN interface. Range: 1-4094.
	port-channel <number>	LAG interface. Range is 1-4094.
Default	BGP uses best local address	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.4070	Added IPv6 and IPv4 support
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.2 update-source vlan 10`	
Related Commands	N/A	
Note	If BGP update-source on neighbor is configured, the given interface's primary address is used as the source address. If BGP update-source configured on a peer group, the primary address is not guaranteed to be the source.	

neighbor no-update-source

	neighbor <ip-address> no-update-source Disables configured source-address for routing updates and for TCP connection establishment with a peer and prevents the system from inheriting this parameter's group configuration.
Syntax Description	N/A
Default	BGP uses best local address
Configuration Mode	config router bgp
History	3.6.3004
Role	admin
Example	`switch (config router bgp 100)# neighbor 10.10.10.2 no-update-source`
Related Commands	N/A
Note	

neighbor weight

neighbor {<ip-address> | <peer-group-name>} weight <value>
no neighbor {<ip-address> | <peer-group-name>} weight

Assigns a weight attribute to paths from the specified neighbor.
The no form of the command resets to default values.

Syntax Description	ipv4_addr, ipv6_addr	IP address of the neighbor
	peer-group-name	Peer group name
	value	Weight value • Range: 1-65535 • Explicitly configuring a default value prevents the system from inheriting this parameter's group configuration
Default	Value is 32768 for router-originated paths and 0 for routes received through BGP	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added IPv6 and IPv4 support
Role	admin	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 weight 100`	
Related Commands	N/A	
Note	• Weight values set through route map commands have precedence over neighbor weight command values. • Other attributes are used only when all paths to the prefix have the same weight. • A path's BGP weight is also configurable through route maps. • When multiple paths to a destination prefix exist, the best-path selection algorithm prefers the path with the highest weight. • Weight is the first parameter that the BGP best-path selection algorithm considers.	

network

network <ip_prefix length> [<route-map-name>]
no network <ip_prefix length> [<route-map-name>]

Configures a route for advertisement to BGP peers.
The no form of the command removes the route from the BGP routes table, preventing its advertisement. The route is only advertised if the router has a gateway to the destination.

Syntax Description	ip_prefix_length	A string that specific route map is assigned to the network.
	length	/24 or 255.255.255.0 format.
	route-map-name	The name of a route-map which is used to set the route's attributes when it is advertised.
Default	N/A	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description, syntax description and notes
Role	admin	
Example	`switch (config router bgp 100)# network 10.10.10.0 /24 routemap`	
Related Commands		
Note	• The parameters "ip-prefix" and "length" specify the route destination • The configuration zeros the host portion of the specified network address (e.g. 192.0.2.4/24 is stored as 192.0.2.0/24) • This command cannot be used with route-maps	

redistribute

redistribute {connected | static | ospf | ospf-internal | ospf-external} [<route-map>|
no redistribute {connected | static | ospf}

Enables redistribution of specified routes to the BGP domain.
The no form of the command disables route redistribution from the specified source.

Syntax Description	connected	Redistributes the direct routes
	static	Redistributes the user-defined (static) route
	ospf	Redistributes all routes learned by OSPF protocol
	ospf-internal	Redistributes all OSPF-learned routes which are marked as internal
	ospf-external	Redistributes all OSPF-learned routes which are marked as external
Default	No redistribution	
Configuration Mode	config router bgp	
History	3.2.1000	
Role	admin	
Example	`switch (config router bgp 100)# redistribute ospf`	
Related Commands	N/A	
Note	• Multiple redistribution options can be applied • This command cannot be used with route-maps	

router-id

router-id <ip-address> [force]
no router-id [force]

Configures a fixed router ID for BGP.
The no form of the command removes the fixed router ID and restores the system default.

Syntax Description	ip-address	IP Address identified the router ID
	force	Applies configuration while BGP is admin-up
Default	The Router ID is dynamically elected (no router-id). • If a loopback interface is configured, the router ID is set to the IP address of the loopback interface. • If multiple loopback interfaces are configured, the router ID is set to the IP address of the loopback interface with the highest IP address. • If no loopback interface is configured, the router ID is set to the highest IP address on a physical interface.	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.3004	Added “force” parameter
Role	admin	
Example	`switch (config router bgp 100)# router-id 10.10.10.10`	
Related Commands		
Note	The IP address configured identifies the BGP speaker. The command triggers an automatic notification and session reset for the BGP neighbors.	

network

[no] network <ip_prefix length> [<route_map_name>]

Adds the given prefix to advertisements that are sent with the specified address family.
The no form of the command removes the given prefix from advertisements that are sent with the specified address family.

Syntax Description	ipv4_prefix	Subnet IP
	route_map_name	Route map name that modifies default route attributes
	length	Allowed prefix length is 24 for IPv4 and IPv6 or 255.255.0.0 for IPv4
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 100)# network 10.10.10.0 /24 default`	
Related Commands		
Note	Address family is identified by the network address itself and not by the configuration command context	

redistribute

[no] [neighbor <peer_group>] redistribute {connected | static}

Enables redistribution of routes to BGP in the specified address family or a peer-group in the address family.
The no form of the command disables the redistribution of routes to BGP.

Syntax Description	connected	Redistributes direct routes
	static	Redistributes static routes
	peer_group	Route map name that modifies default route attributes
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 100)# redistribute connected`	
Related Commands		
Note		

route-map

[no] [neighbor <peer_group>] route-map <route_map_name> [{in | out}]

Specifies a route map that will be applied in the given direction for specific address family.

Syntax Description	route_map_name	Name of a route map to apply.
	in/out	Specifies in which direction the route map is applied. If nothing is given - route map is applied in both directions
	peer_group	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
Role	admin	
Example	`switch (config router bgp 100)# route-map default in`	
Related Commands		
Note		

6.4.8.3 Show

show {ip | ipv6} bgp

show {ip | ipv6} bgp [vrf <vrf-name>] [<ipv4-prefix> <length> [detail | longer-prefixes [detail]]]

Displays information about the BGP routes table (RIB).

Syntax Description	ipv4_prefix, ipv6_prefix	IPv4 and IPv6 subnet
	length	Netmask (e.g. /24 or 255.255.255.0).
	detail	Displays detailed information about a subset of the bgp learned routes.
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination. Example: If "10.20.30.0 /24 longer-prefixes" is run, all routes starting with 10.20.30 regardless of the prefix length (10.20.30.X /24, 10.20.30.X /25, etc.) are displayed – providing there are any such routes received/sent from/to that neighbor.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6
	3.6.6000	Updated Example for "detail" parameter
	3.7.100x	Updated example "show ip bgp" Updated example for "longer-prefixes" parameter
Role	admin	

Example

```
switch (config) # show ip bgp 192.168.100.0 /24

BGP table version: 22
Local router ID: 192.168.100.11

Status codes:
  s: suppressed
  d: damped
  h: history
  *: valid
  >: best
  i: internal
  r: RIB-failure
  S: Stale
  m: multipath
  b: backup-path
  x: best-external

Origin codes:
  i: IGP
  e: EGP
  ?: incomplete

------------------------------------------------------------------------
Network           Next Hop  Status  Metric   LocPrf   Weight   Path
------------------------------------------------------------------------
192.168.100.0/24   0.0.0.0    *>     0       100      32768     i
```

Related Commands

N/A

Note

Aggregated information in the “detail” parameter (i.e. aggregator AS, aggregator ID) is displayed only for aggregated routes.
Generic and “Longer prefixes” examples were updated.

show ip bgp address-family

show ip bgp address-family [vrf <vrf-name>] {l2vpn-evpn | <ipv4-unicast | ipv6-unicast>} [active] [detail]

Displays address-family configuration.

Syntax Description	l2vpn-evpn	Displays information about L2VPN-EVPN address family.
	active	Displays active neighbors in that address family (configured, active or dynamic)
	detail	Displays detailed info about configuration and configured/active neighbors for the specified address-family
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4070	
	3.7.100x	Added “l2vpn-evpn” parameter and updated Example
Role	admin	
Example		

```
switch (config) # show ip bgp address-family l2vpn-evpn

Address family               : L2VPN EVPN
Maximum Path                 : 0/0
Redistribute                 :
Total Neighbors              : 1
Total peer-groups            : 1
Total dynamic ranges         : 0
switch (config) # show ip bgp address-family l2vpn-evpn active
Address family               : L2VPN EVPN
Networks                     :
maximum-path                 : 0/0
redistribute                 : -
Total neighbors              : 2
Total peer-groups            : 0
Total dynamic ranges         : 0
switch (config) # show ip bgp address-family l2vpn-evpn detail

Address family    : L2VPN EVPN
Maximum Path      : 0/0
Redistribute      :
Total Neighbors   : 1

Neighbors:
  ------------------------------------------------------------------------------------
  Neighbor   V   AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down   State/PfxRcd
  ------------------------------------------------------------------------------------
  1.1.1.1    4   65002   0         1         6        0     0      Never     ACTIVE/0

Total peer-groups         : 1
Peer Group                : peer
Total dynamic ranges      : 0
```

Related Commands N/A

Note

show ip bgp community

	show ip bgp [vrf <vrf-name>] community <comm$_1$> <comm$_2$> … <comm$_n$> [exact] [detail] Displays information about the BGP routes (RIB) filtered according to communities.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.0000
Role	admin
Example	
Related Commands	show ip bgp
Note	

show ip bgp evpn

show ip bgp [vrf <vrf-name>] [neighbors <ip | peer-group> [received | advertised]] evpn [route-type <type> | community {<aa:nn> | <number>} | extcommunity route-target {<aa:id> | <aa.bb:id> | <ip:id>} | extcommunity router-mac <mac-address> | vni <value> | rd *] [detail]

Displays BGP EVPN routes received from all neighbors in specified VRF or the VRF currently under context.

Syntax Description	ipv4_addr	Neighbor IP address
	peer_group	Peer group name
	route-type	Possible values: 1-5 • 1 – Ethernet Auto-discovery Route • 2 – MAC/IP Advertisement Route • 3 – Inclusive Multicast Ethernet Tag Route • 4 – Ethernet Segment Route • 5 – IP Prefix Route
	community	<aa:nn> – community number <number> – community number
	extcommunity route-target	Filters by route target <aa:id> – Route Target (asplain) <aa.bb:id> – Route Target (asdot) <ip:id> – Rout Target (IP)
	extcommunity router-mac	Filters by router MAC
	vni	VNI value. Range: 1-16777215.
	rd	Filters by route target <aa:id> – Route Target (asplain) <aa.bb:id> – Route Target (asdot) <ip:id> – Rout Target (IP)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8100	
Role	admin	
Example		

```
switch (config) # show ip bgp evpn

-------------------------------------------------------------------------------------------------------------------
RD                Type            Data                            Next Hop      Metric   LocPrf   Weight   Path
-------------------------------------------------------------------------------------------------------------------
192.168.3.2:65001   auto-discovery   00:00:00:00:00:00:00:00:00:00    192.168.3.2   0        100      0        65002 ?
0.0.0.0:0           mac-ip           00:01:02:03:04:08 192.168.1.1    192.168.3.2   0        100      0        65002 ?
0.0.0.0:0           mac-ip           00:01:02:03:04:08 192.168.10.1   192.168.3.2   0        100      0        65002 ?
253.233.0.0:10      mac-ip           00:01:02:03:04:07 192.168.4.2    192.168.5.2   0        100      300      65003 ?

switch (config) # show ip bgp evpn detail

1 paths for auto-discovery 00:00:00:00:00:00:00:00:00:00 Route Distinguisher:192.168.3.2:65001:
  65002:
    next hop                          : 192.168.3.2
    neighbor ip                       : 192.168.3.2
    router id                         : 192.168.3.2
    metric                            : 0
    weight                            : 0
    local pref                        : 100
    origin                            : incomplete
    Extended Community                : 0:0(Route-Target-AS)
    flags                             : valid, best
    vni                               : 200

1 paths for mac-ip 00:01:02:03:04:08 192.168.1.1 Route Distinguisher:0.0.0.0:0:
  65002:
    next hop                          : 192.168.3.2
    neighbor ip                       : 192.168.3.2
    router id                         : 192.168.3.2
    metric                            : 0
    weight                            : 0
    local pref                        : 100
    origin                            : incomplete
    Extended Community                : 100:10(Route-Target-AS)
    Extended Community                : tunnelTypeVxlan(TunnelEncap)
    flags                             : valid, best
    esi                               : 00:00:00:00:00:00:00:00:00:00
    vni                               : 1
```

Related Commands

Note

show ip bgp evpn summary

show ip bgp [vrf <vrf>] evpn summary

Displays some basic statistics about BGP per VRF only for neighbors who support L2EVPN AF.

Syntax Description	vrf	Name of VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8100	
Role	admin	

Example

```
switch (config) # show ip bgp evpn summary

VRF name                                : vrf-default
BGP router identifier                   : 192.168.5.1
local AS number                         : 65001
BGP table version                       : 2
Main routing table version              : 2
IPV4 Prefixes                           : 0
IPV6 Prefixes                           : 0
L2VPN EVPN Prefixes                     : 1

-------------------------------------------------------------------------------------------------
Neighbor      V   AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down      State/PfxRcd
-------------------------------------------------------------------------------------------------
192.168.3.2   4   65002   25        29        2        0     0      0:00:11:10   ESTABLISHED/1
192.168.5.2   4   65003   24        28        2        0     0      0:00:11:17   ESTABLISHED/0
```

Related Commands

Note

show ip bgp neighbors

show {<ip>} bgp neighbors [vrf <vrf-name>] [<ip-address>]

Displays summaries information about all BGP neighbors.

Syntax Description	ip-prefix	Destination to aggregate
	length	Mask length (e.g. /24 or 255.255.255.254)
	ip-address	neighbor address
	longer-prefixes	Displays information about routes with longer prefixes than given
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.100x	Updated Example
Role	admin	

Example

```
switch (config) # show ip bgp neighbors 192.168.2.2

BGP neighbor: 192.168.2.2, remote AS: 65002, link: external:
BGP version                                 : 4
  Configured hold time in seconds           : 180
  keepalive interval in seconds             : 60
  Minimum holdtime from neighbor in seconds: 90
  Peer group                                :

  Neighbor configuration:
    ------------------------------------------------------------------------
    Configuration                   IPV4 Unicast  IPV6 Unicast  L2VPN EVPN
    ------------------------------------------------------------------------
    Configured AFI SAFI             Enabled       Disabled      Enabled
    Send Community                  Disabled      Disabled      Disabled
    Send Extended Community         Disabled      Disabled      Disabled
    Route Reflection                Disabled      Disabled      Disabled
    Next Hop Unchanged              Disabled      Disabled      Disabled

  Neighbor capabilities:
    Route Refresh              : advertise and received
    Soft Reconfiguration       : Disabled
    Graceful Restart Capability: advertise
    Address family IPv4 Unicast: advertise and received
    Address family IPv6 Unicast: n/a
    Address family L2VPN EVPN  : advertise and received
```

```
  Message statistics:
    InQ depth : 0
    OutQ depth: 0

    ----------------------------------------
    Parameter             Sent         Rcvd
    ----------------------------------------
    Opens                 1            1
    Notification          0            0
    Updates               3            2
    Keepalives            12           11
    Refreshes             0            0
    Total                 16           14

    Default minimum time between advertisement runs in seconds: 30

  L2VPN EVPN:
    ----------------------------------------------
    Prefix activity       Sent         Rcvd
    ----------------------------------------------
    Prefixes Current      2            2
    Prefixes Total        2            2
    Implicit Withdraw     0            0
    Explicit Withdraw     0            0
    Used as bestpath      n/a          2
    Used as multipath     n/a          n/a

    ------------------------------------------------------
    Local Policy Denied Prefixes    Outbound     Inbound
    ------------------------------------------------------
    Total                           0            0

  Connection Information:
    Connections established              : 4
    Dropped                              : 1
    Last Reset                           : 0:00:03:22
    Last Drop Reason                     : 6 (2)
    Maximum hops to external BGP neighbor: 255
    Connection State                     : ESTABLISHED
    Local host                           : 192.168.2.1
    Local port                           : 179
    Foreign host                         : 192.168.2.2
    Local Port                           : 50394


switch (config) # show ip bgp neighbors

BGP neighbor: 192.168.2.2, remote AS: 65001, link: internal:
  BGP version                             : 4
  Configured hold time in seconds         : 180
  keepalive interval in seconds           : 60
  Minimum holdtime from neighbor in seconds: 90
```

Related Commands N/A

Note

show ip bgp neighbors received

show ip bgp neighbors <ip-address> received [<ip-address> [<mask>] [longer-prefixes]]

Displays BGP summary information.

Syntax Description	ip-address	Neighbor IP address
	mask	Mask length
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination. (Only available if both IP and mask are specified.)
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.100x	
Role	admin	
Example		
Related Commands	N/A	
Note		

show ip bgp neighbors received detail

show ip bgp neighbors <ip-address> received [<ip-address> [<mask> [longer-prefixes]]] detail

Displays detailed information on routes received from neighbors.

Syntax Description	ip-address	Neighbor IP address. Provide optionally to display routes received from specified neighbor.
	mask	Mask length. Displays routes received from specified neighbor filtered by the specified network.
	longer-prefixes	Displays routes received from specified neighbor filtered by the specified prefix and longer
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.100x	Updated output
Role	admin	

Example

```
switch (config)# show ip bgp 192.168.100.0 /24 longer-prefixes detail

BGP routing table entry for: 192.168.100.0/24
Version                    : 22
Paths                      : (1, best: #1)

Local Connected:
  Origin    : IGP
  metric    : 0
  localpref : 100
  weight    : 32768
  Attributes: valid, best
switch (config)# show ip bgp 192.168.100.0 /24 detail

BGP routing table entry for: 192.168.100.0/24
Version                    : 22
Paths                      : (1, best: #1)

Local connected:
  0.0.0.0 from 0.0.0.0 (192.168.100.11):
    Origin    : IGP
    metric    : 0
    localpref : 100
    weight    : 32768
    Attributes: valid, sourced, best
```

Related Commands	N/A
Note	

show ip bgp paths

show ip bgp paths [vrf <vrf-name>] [ipv4 | ipv6]

Displays summary of all AS paths and for prefixes for specific address family.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.5200 3.6.4070 Added support for IPv4 and IPv6
Role	admin

Example	`switch (config) # show ip bgp paths`
	`Refcount  Metric  Path`
	`1  0  4 50 100`
	`1  0  2 50 100`
	`1  0  4 40`
	`1  0  12 50 100`
	`1  0  2`
	`1  0  2 20`
	`switch (config) #`

Related Commands	N/A
Note	

show ip bgp peer-group

show ip bgp peer-group [vrf <vrf-name>] [peer-group-name] [address-family <ip-address>]

Displays information about peer groups and configuration, filtered per address family.

Syntax Description	peer-group-name	Displays information about a specific peer-group.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
	3.6.8100	Updated Example
	3.7.100x	Updated Example
Role	admin	

Example

```
Name                    : peerGrp1
Hold time               : 180
Keep-alive              : 60
Max prefix              : 100000
Weight                  : 0
Export local preferences: 100
Import local preferences: 100
Status Down             : no
EBGP Multihop           : 1
Next Hop Self           : no
Soft Reconfiguration    : no
Next Hop Peer           : no
Remove Private AS       : no
Transport Mode          : no
Password                : no
Local AS                : 0
No Prepend              : no
Replace AS              : no
Soft Reconfiguration    : Disabled

-------------------------------------------------------------------------
Configuration                  IPV4 Unicast  IPV6 Unicast  L2VPN EVPN
-------------------------------------------------------------------------
Configured AFI SAFI            Disabled      Disabled      Disabled
Send Community                 Disabled      Disabled      Disabled
Send Extended Community        Disabled      Disabled      Disabled
Route Reflection               Disabled      Disabled      Disabled
Next Hop Unchanged             Disabled      Disabled      Disabled

-----------------------------------------------------------------------------------------------
Neighbor        V    AS     MsgRcvd  MsgSent  TblVer   InQ   OutQ  Up/Down      State/PfxRcd
-----------------------------------------------------------------------------------------------
192.168.2.2     4    65001  355      413      7        0     0     0:00:00:26    ESTABLISHED/2
```

Related Commands	N/A
Note	

show ip bgp summary

show ipv6 bgp {<id> | all} summary [vrf <vrf-name>]

Displays BGP summary for IPv6 addresses.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.5200 3.6.4070 Added support for IPv6
Role	admin

Example	`switch (config) # show ip bgp summary`
	`BGP router identifier 3.5.7.4, local AS number 4`
	`BGP table version is 70/120, main routing table version 70/96`
	`BGP using 26308 total bytes of memory`
	`BGP activity 37/8 IPv4 prefixes, 37/8 IPv6 prefixes, 37/4 paths`
	`Neighbor V AS MsgRcvd MsgSent InQ OutQ Up/Down State/PfxRcd`
	`2001::1 4 7 3 9 0 0 0:00:00:48 ESTABLISHED/total num-ber of prefixes`

Related Commands	N/A
Note	

show ip bgp update-group

show ip bgp update-group [<neighbor ip address>]

Displays update-group information for all neighbors.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4070 3.7.100x Updated Example
Role	admin
Example	

```
r-mgtswd-270 [standalone: master] (config) # show ip bgp update-group 192.168.2.2

Update-group for neighbor: 192.168.2.2
BGP router identifier    : 192.168.2.1
local AS number          : 65001
BGP table version        : 7


------------------------------------------------------------------------------------------------
Neighbor        V   AS      MsgRcvd  MsgSent  TblVer   InQ   OutQ  Up/Down     State/PfxRcd
------------------------------------------------------------------------------------------------
192.168.2.2     4   65001    368      428       7        0     0     0:00:06:30   ESTABLISHED/2


r-mgtswd-270 [standalone: master] (config) # show ip bgp update-group

Update-group                                    : 5
BGP version                                     : 4
Address Family                                  : IPv4 Unicast
Minimum time between advertisements runs in seconds: 30

Has 1 members:
  192.168.2.2

Update-group                                    : 6
BGP version                                     : 4
Address Family                                  : L2VPN EVPN
Minimum time between advertisements runs in seconds: 30

Has 1 members:
  192.168.2.2
```

Related Commands	N/A
Note	

show ip bgp vrf summary

show ip bgp vrf {<vrf-name> | all} summary

Displays BGP summary info for all or specified VRFs.

Syntax Description	vrf-name	Displays BGP summary for specified VRF
	all	Displays BGP summary for all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.8100	Updated Example
Role	admin	

Example

```
switch (config)# show ip bgp summary

VRF name                                : vrf-default
BGP router identifier                   : 1.1.1.2
local AS number                         : 65001
BGP table version                       : 3
Main routing table version              : 3
IPV4 Prefixes                           : 0
IPV6 Prefixes                           : 0
L2VPN EVPN Prefixes                     : 2

-------------------------------------------------------------------------------------------
Neighbor   V    AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down      State/PfxRcd
-------------------------------------------------------------------------------------------
1.1.1.1    4    65002   25        29        3        0     0      0:00:10:38   ESTABLISHED/2
1.1.1.5    4    100     0         0         3        0     0      Never        IDLE/0
```

Related Commands

Note

6.4.8.4 IP AS-Path Access-List

ip as-path access-list

ip as-path access-list <list-name> {permit | deny} <reg-exp> [any | egp | igp | incomplete]
no ip as-path access-list <list-name>

Creates an access list to filter BGP route updates.
The no ip as-path access-list command deletes the named access list.

Syntax Description	list-name	The name for the access list
	permit	Permits access for a matching condition
	deny	Denies access for a matching condition
	reg-exp	Regular expression that is used to specify a pattern to match against an input string.
	any	Any route type
	egp	External BGP routes
	igp	Internal BGP routes
	incomplete	Routes marked as “Incomplete”
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# ip as-path access-list mylist permit` `switch (config)#`	
Related Commands	N/A	
Note	If access list_name does not exist, this command creates it. If it already exists, this command appends statements to the list.	

show ip as-path access-list

show ip as-path access-list [list-name]

Presents defined as-path access lists

Syntax Description	list-name	Displays a specific prefix-list.
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# show ip as-path access-list mylist`	
Related Commands	N/A	
Note		

6.4.8.5 IP Community-List

ip community-list standard

ip community-list standard <list-name> {deny | permit} <list-of-communities>
no ip community-list standard <list-name>

Adds a standard entry to a community-list.
The no form of the command deletes the specified community list.

Syntax Description	list-name	The name for the community list
	permit	Permits access for a matching condition.
	deny	Denies access for a matching condition.
	list-of-communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# ip community-list standard mycommunity permit 1:2 3:4`	
Related Commands	N/A	
Note	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

ip community-list expanded

ip community-list expanded <list-name> {deny | permit} <reg-exp>
no ip community-list expanded <list-name>

Adds a regular expression entry to a community-list
The no form of the command deletes the specified community list.

Syntax Description	list-name	Configures a named standard community list.
	permit	Permits access for a matching condition.
	deny	Denies access for a matching condition.
	reg-exp	Regular expression that is used to specify a pattern to match against an input string.
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# ip community-list expanded mycommunity permit 1:[0-9]+`	
Related Commands	N/A	
Note	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

show ip community-list

show ip community-list [community-list-name]

Displays the defined community lists

Syntax Description	community-list-name	An optional parameter to display only the specified list
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Role	admin	
Example	`switch (config)# show ip community-list mycommunity`	
Related Commands	N/A	
Note	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

6.5 BFD Infrastructure

Many protocols uses slow Hello mechanisms and failure is detection usually seconds after the problem occurs. The BFD goal is to provide low overhead short duration detection of failures between adjacent nodes and single mechanism that can be used for liveness detection over any media.

BFD session is established by the application that uses it. There is no discovery mechanism. E.g. in OSPF BFD session is established to neighbors that were discovered by OSPF hello protocol.

BFD supports multiple modes: one of them is Asynchronous.

In Asynchronous mode a system periodically sends BFD packets to verify connectivity. If a number of packets in a row are not received – the session is declared down.

A system can be passive or active. Active system initiates BDF sessions. Both systems can be active. (Only active mode is supported.)

6.5.1 Session Establishment

A session begins with exchange of control packets. When bidirectional communication is achieved - a session becomes Up.

After session becomes up - control packet rate can be incremented.

Each side informs the neighbor in what intervals it is going to send BFD packets and what minimum interval it can receive BFD packets is.

Detection time is different in both directions and depends on negotiated parameters.

In Asynchronous mode—agreed transmit interval or remote system—max between local minimum rx time and last received min transmit time.

Detection time is equal to agreed transmit interval of remote system multiplied to multiplier received from remote system.

6.5.2 Interaction with Protocols

BFD session can be single-hop or multi-hop:

- Single hop session traverse between two adjacent IP neighbors. BFD control packet should be encapsulated in UDP with DPORT = 3784. SPORT should be in range 49152 to 65535. Same SPORT must be used for all control BFD packets for given session and is unique between different sessions. TTL value is 255.
- Multi-hop sessions traverse between to remote ip neighbors. Control packets are encapsulated in UDP with DPORT = 4784.

If different protocols want to establish a BFD session with the same remote system for same data plane - they should share BFD session.

IPv4 and IPv6 data protocols have different BFD sessions.

In OSPF Protocol neighbor discovery protocol establishes single hop BFD sessions. For OSPF when session fails - it tears down OSPF neighbor.

BFD session is established to BGP neighbor (single hop or multiple hop).

Single hop BFD session can be established for static route next hop.

6.5.3 Config Commands

protocol bfd

	[no] protocol bfd Enables bfd on a system level The no form of the command removes bfd configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config router bgp
History	3.6.4070
Role	admin
Example	`switch (config router bgp)# bfd shutdown` `switch (config router bgp)#`
Related Commands	
Note	The command will return an error if BFD is enabled in clients already running on the system (static routes or BGP of OSPF).

bfd shutdown

	[no] bfd shutdown [vrf <vrf-name>] Disables bfd sessions but doesn't remove the configuration if VRF is not given the command will be executed in active VRF.
Syntax Description	N/A
Default	N/A
Configuration Mode	config router bgp
History	3.6.4070
Role	admin
Example	`switch (config router bgp)# bfd shutdown` `switch (config router bgp)#`
Related Commands	
Note	"no ip bfd shutdown" or BFD interval parameters modification are affect traffic for all protocols; OSPF, BGP, static routes. The dynamic protocols (OSPF and BGP) will restore the connection based on Hello protocol. For static routes, please execute "no ip route static bfd <ip address>"

6.5.4 Interface Commands

bfd interval

bfd interval [vrf <vrf-name>] [transmit_rate] [min_rx] [multiplier]
no bfd interval

Sets the interval rates between BFD messages.
The no form of the command removes bfd interval rates.

Syntax Description	transmit_rate	Transfer time between two consecutive BFD messages, the actual time is negotiated between two systems. Range: 50-60000 (msec) Default: 300 (msec)
	min_rx	Minimum time between neighbor messages, the actual time is negotiated between two systems. Range: 50-60000 (msec) Default: 150 (msec)
	multiplier	Defines a time period to detect BFD failure. Range: 3-50 Default: 3
Default	N/A	
Configuration Mode	config	
History	3.7.11xx	Updated example
	3.6.4070	
Role	admin	
Example	`switch (config router bgp)# ip bfd interval transmit-rate 300 multi-plier 3 min-rx 300 force`	
Related Commands		
Note	The command is executed in Active VRF when VRF is not specified.	

ip ospf bfd

	[no] ip ospf bfd Enables BFD on the given interface for all OSPF neighbors on a number of active sessions. The no form of the command disables BFD on all OSPF neighbors.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ethernet
History	3.6.4070 3.6.4110 Added “no” form of the command.
Role	admin
Example	`switch (config )# ip ospf bfd`
Related Commands	N/A
Note	"ip ospf bfd" affects traffic, OSPF will restore the connection based on Hello protocol.

ip route bfd

ip route [vrf <vrf_name>] <prefix> <next_hop> bfd
[no] ip route [vrf <vrf_name>] <prefix> <next_hop> bfd

Configures static route with BFD enabled on a specified VRF.
The no form of the commands removes the route.

Syntax Description	vrf-name	VRF session name
	prefix	Subnet IP address
	next_hop	ip address of next hop
Default	N/A	
Configuration Mode	config	
History	3.6.4070	
	3.7.11xx	Updated command, syntax and example
Role	admin	
Example	`switch (config) # ip route vrf default 1.1.1.0/24 3.3.3.3 bfd`	
Related Commands	N/A	
Notes	When a session fails, all static routes pointing to the specified gateway are removed from routing decision.	

show ip route static

show ip route [vrf [<vrf-name> | all]] static

Displays static routing table of VRF instance.

Syntax Description	all	Displays routing tables for all VRF instances
	vrf	vrf name
Default	Default vrf	
Configuration Mode	Any command mode	
History	3.7.11xx	Update command syntax
	3.6.4070	
Role	admin	
Example	`switch (config) # show ip route vrf default static`	
Related Commands	ip route	
Notes	• If no routing-context is specified, the “routing-context” VRF is automatically displayed	

show ip bfd neighbors

show ip bfd [vrf <name>|all] neighbors [brief|<ip>]

Displays bfd table of neighbor VRF instances.

Syntax Description	all	Displays tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Role	admin	

Example

```
switch (config) # show ip bgp neighbors 1000::1040
BGP neighbor: 1000::1040, remote AS: 100, link: external
  BGP version: 4, remote router ID: 2.1.1.1
  BGP State: ESTABLISHED
  Last read: 0:00:09:28, last write: 0:00:09:28, hold time is: 180, keepalive interval in seconds: 60
  BFD State: Up
  Configured hold time in seconds: 180, keepalive interval in seconds: 60
  Minimum holdtime from neighbor in seconds: 180

  Neighbor capabilities:
Route refresh: advertise and received
Graceful Restart Capability: advertise and received
Address family IPv4 Unicast: advertise and received
Address family IPv6 Unicast: n/a

  Message statistics:
InQ depth is: 0
OutQ depth is: 0
                      ----   -----
                    Sent   Rcvd
                      ----   -----
Opens:              1      1
Notifications:      0      0
Updates:            4      4
Keepalives:      1587   1593
Route Refresh:      0      0
Total:           1592   1598
  Default minimum time between advertisement runs in seconds: 30

  For address family: IPv4 Unicast
  BGP table version: 7
  Output queue size : 0

                             ----   ----
                            Sent   Rcvd
  Prefix activity:          ----   ----
Prefixes Current:           4      2
Prefixes Total:             4      2
Implicit Withdraw:          0      0
Explicit Withdraw:          0      0
Used as bestpath:         n/a      2
Used as multipath:        n/a    n/a

                                  ---------------
                               OutboundInbound
  Local Policy Denied Prefixes:---------------
Total:                            0      0

  Connections established: 1; dropped: 1
  Last reset: 0:23:01:17, due to: 0 (0)
External BGP neighbor possible distance in hops: 1
Connection state is: ESTABLISHED
Local host: 1.1.1.1, Local port: 49616
Foreign host: 1000::1040, Foreign port: 179
```

Related Commands

Notes

6.6 Policy Rules

6.6.1 Route Map

Route maps define conditions for redistributing routes between routing protocols. A route map clause is identified by a name, filter type (permit or deny) and a sequence number. Clauses with the same name are components of a single route map; the sequence number determines the order in which the clauses are compared to a route.

Route maps can be used only for the BGP protocol.

Route maps cannot be used for the commands "network" on page 1429 or "redistribute" on page 1430.

6.6.1.1 Commands

route-map

route-map <map-name> [deny | permit] [sequence-number]
no route-map <map-tag> {deny | permit} [<sequence-number>]

Creates a route map that can be used for importing, exporting routes and applying local policies.

Syntax Description	name	Name of the route-map
	deny | permit	Configures the rule to be used
	sequence-number	Sequence number for a route-map specific record
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
	3.3.5200	Updated notes
Role	admin	
Example	`switch (config) # route-map mymap permit 1200` `switch (config route-map mymap permit 1200)#`	
Related Commands	N/A	
Note	• All changes in a the route map configuration mode become pending until the end of the route-map session. • If not configured, deny | permit is configured as permit • If not configured, sequence-number default value is 10	

continue <sequence-number>

	continue <sequence-number> **no continue** Enables additional route map evaluation of routes whose parameters meet the clause's matching criteria. The no form of the command removes this configuration from the route map clause.
Syntax Description	sequence-number
Default	N/A
Configuration Mode	config route map
History	3.3.5006 3.3.5200 Updated example
Role	admin

Example	`switch (config route-map mymap permit 10)# match as-number 40`
	`switch (config route-map mymap permit 10)# set weight 7`
	`switch (config route-map mymap permit 10)# continue 1200`
	`switch (config route-map mymap permit 10)# exit`
	`switch (config)# show route-map test`
	`route-map test, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`    continue 1200`
	`switch (config route-map mymap permit 10)# route-map test permit 10 no continue`
	`switch (config route-map mymap permit 10)# show route-map test`
	`route-map test, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`switch (config route-map mymap permit 10)#`

Related Commands	route-map <map-name> [deny | permit] [sequence-number]
Note	• A clause typically contains a match (route-map) and a set (route-map) statement. The evaluation of routes whose settings are the same as match statement parameters normally end and the clause's set statement are applied to the route. Routes that match a clause containing a continue statement are evaluated against the clause specified by the continue statement. • When a route matches multiple route-map clauses, the filter action (deny or permit) is determined by the last clause that the route matches. The set statements in all clauses matching the route are applied to the route after the route map evaluation is complete. Multiple set statements are applied in the same order by which the route was evaluated against the clauses containing them. • Continue cannot be set to go back to a previous clause; <sequence-number> of the continue must always be higher than the current clause's sequence number.

abort

	abort Discards pending changes and returns to global configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config route map
History	3.3.5006 3.3.5200 Updated example
Role	admin

Example	`switch (config)# route-map mymap permit 10 match as-number 40`
	`switch (config)# route-map mymap permit 10 set weight 7`
	`switch (config)# show route-map test`
	`route-map test, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`switch (config)# route-map mymap permit 1200`
	`switch (config route-map mymap permit 1200)# set weight 11`
	`switch (config route-map mymap permit 1200)# abort`
	`switch (config)# show route-map mymap`
	`route-map mymap, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`switch (config)#`

Related Commands	N/A
Note	

exit

	exit Saves pending route map clause changes to running-config and returns to global configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config route map
History	3.3.5006
Role	admin

Example	`switch (config)# route-map mymap permit 10 match as-number 40`
	`switch (config)# route-map mymap permit 10 set weight 7`
	`switch (config)# show route-map test`
	`route-map test, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`switch (config)# route-map mymap permit 1200`
	`switch (config route-map mymap permit 1200)# set weight 11`
	`switch (config route-map mymap permit 1200)# exit`
	`switch (config)# show route-map test`
	`route-map mymap, permit, sequence 10`
	`  Match clauses:`
	`    as-number 40`
	`  Set clauses:`
	`    weight 7`
	`route-map mymap, permit, sequence 1200`
	`  Set clauses:`
	`    weight 11`
	`switch (config)#`

Related Commands	N/A
Note	

match as-number

match as-number <number>
no match as-number

Filters according to one of the AS numbers in the AS path of the route.
The no form of the command removes this configuration from the route map clause.

Syntax Description	number	Autonomous system number to check.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Role	admin	
Example	`switch (config route-map mymap permit 10)# match as-number 40` `switch (config route-map mymap permit 10)#`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

match as-path

match as-path <as-path-list name>
no match as-path

Creates a route map clause entry that matches the route‘s AS path using an as-path access-list.
The no form of the command removes the match statement from the configuration mode route map clause.

Syntax Description	number	Autonomous system number to check.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
	3.6.3004	Added note
Role	admin	
Example	`switch (config route-map mymap permit 10)# match as-path my-list`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route’s properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied • An as-path-list must already exist before a node is configured to use it	

match community

match community <list-of-communities> [exact-match]
no match community <list-of-communities>

Creates a route map clause entry that matches a route if it contains at least the specified communities.
The no form of the command removes the match clause.

Syntax Description	list of communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
	exact-match	Creates a route map clause entry that matches the route‘s communities exactly.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Role	admin	
Example	`switch (config route-map mymap permit 10)# match community 1:100 3:52`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route’s properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied. • Route-map’s match on a list of communities is performed with the command “match community-list” and not this command.	

match community-list

match community <communities-list-name> exact-match
no match community <communities-list-name> exact-match

Creates a route map clause entry that specifies one route filtering condition
The no form of the command removes the match clause.

Syntax Description	communities-list-name A name of an IP community list
Default	N/A
Configuration Mode	config route map
History	3.3.5006
Role	admin
Example	`switch (config route-map mymap permit 10)# match community-list COM_LIST exact-match`
Related Commands	N/A
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.

match interface

match interface <interface-type> <number>
no match interface

Matches the route's interface
The no form of the command removes the match clause.

Syntax Description	prefix-list-name	Prefix-list name.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Role	admin	
Example	`switch (config route-map mymap permit 10)# match  1/1`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

match ip address

match ip address <prefix-list-name>
no match ip address

Filters according to IPv4 prefix list.
The no form of the command removes this configuration from the route map clause.

Syntax Description	prefix-list-name	Prefix-list name.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Role	admin	
Example	`switch (config route-map mymap permit 10)# match ip address listSmallRoutes`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied. • The prefix-list-name should point to an existing IP prefix-list. If it is not found, no route is considered as a match for this clause.	

match ip next-hop

match ip next-hop <ipv4/ipv6>
no match ip next-hop

Configures a route's entry next-hop match.
The no form of the command removes a route-map's entry next-hop match.

Syntax Description	ipv4/ipv6	Next hop IP address: A.B.C.D (e.g. 10.0.13.86).
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6
Role	admin	
Example	`switch (config route-map mymap permit 10)# match ip next-hop 10.10.10.10`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

match local-preference

match local-preference <value>
no match local-preference

Configures a route's entry local-preference match.
The no form of the command removes a route-map's entry local-preference match.

Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.4.0000	Updated value range
Role	admin	
Example	`switch (config route-map mymap permit 10)# match local-preference 10`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

match metric

match metric <value>
no match metric

Configures a route's entry metric match.
The no form of the command removes a route-map's entry metric match.

Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.4.0000	Updated value range
Role	admin	
Example	`switch (config route-map mymap permit 10)# match metric 10`	
Related Commands	N/A	
Note	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

set as-path prepend

set as-path prepend <value$_1$> <value$_2$> ... <value$_n$>
no set as-path prepend

Modifies as-path on affected routes
The no form of the command removes the set statement from the route map.

Syntax Description	value	BGP AS number that is prepended to as-path. Range: 1-4294967295.
Default	N/A	
Configuration Mode	config route map	
History	3.4.0000	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set as-path prepend 5 10`	
Related Commands	N/A	
Note		

set as-path tag

set as-path tag <value>
no set as-path tag

Configures a route's entry AS-path tag parameter.
The no form of the command removes a route-map's entry AS path tag setting.

Syntax Description	value	Range: 1-2147483648.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set as-path tag 1`	
Related Commands	N/A	
Note		

set community

	set community {none} **[no] set community {none}** Sets the community attribute of a distributed route The no form of the command removes the set statement from the clause.
Syntax Description	N/A
Default	N/A
Configuration Mode	config route map
History	3.7.11xx Updated syntax 3.3.5200
Role	admin
Example	`switch (config route-map mymap permit 10)# set community 1:2 3:4`
Related Commands	N/A
Note	

set community additive

set community <list-of-communities> additive
no set community <list-of-communities> additive

Adds the matching communities
The no form of the command removes the set statement from the clause.

Syntax Description	list-of-communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set community none`	
Related Commands	N/A	
Note		

set community none

	set community none **no set community none** Sets the community attribute of a distributed route to be empty The no form of the command removes the set statement from the clause.
Default	N/A
Configuration Mode	config route map
History	3.3.5200
Role	admin
Example	`switch (config route-map mymap permit 10)# set community none`
Related Commands	N/A
Note	

set community delete

set community <list of communities> delete
no set community <list of communities> delete

Deletes matching communities.
The no form of the command removes the set statement from the clause.

Syntax Description	list of communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch-e07c04 [standalone: master] (config) # route-map test_route_map` `switch-e07c04 [standalone: master] (config route-map test_route_map permit 10) # set community 400:1 delete`	
Related Commands	N/A	
Note		

set community-list

set community-list <community-list-name>
no set community <list of communities>

Configures a named standard community list.
The no form of the command removes the set statement from the clause.

Syntax Description	<community-list-name>	Name of community list
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.6.3004	Added note
Role	admin	
Example	`switch (config route-map mymap permit 10)# set community internet 1:3 additive`	
Related Commands	N/A	
Note	A community-list must already exist before a node is configured to use it	

set community-list additive

set community-list <community-list-name> additive
no set community <list of communities> additive

Adds to existing communities using the communities found in the community list. The no form of the command removes the set statement from the clause.

Syntax Description	<community-list-name> Name of community list
Default	N/A
Configuration Mode	config route map
History	3.3.5200
Role	admin
Example	`switch (config route-map mymap permit 10)# set community-list mycommunity additive`
Related Commands	N/A
Note	

set community-list delete

set community-list <community-list-name> delete
no set community-list

Deletes the matching community list permit entries from the route community list
The no form of the command removes the set statement from the clause.

Syntax Description	community-list-name	Name of community list
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set community-list mycommu-nity delete`	
Related Commands	N/A	
Note		

set ip next-hop

set ip next-hop <ipv4/ipv6>
no set ip next-hop

Configures a route's entry next-hop parameter.
The no form of the command removes a route-map's entry next-hop setting.

Syntax Description	ipv4/ipv6	Route next-hop IP: A.B.C.D (e.g. 10.0.13.86).
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6
Role	admin	
Example	`switch (config route-map mymap permit 10)# set ip next-hop 10.10.10.10`	
Related Commands	N/A	
Note		

set local-preference

set local-preference <value>
no set local-preference

Configures a route's entry local-preference parameter.
The no form of the command removes a route-map's entry local-pref setting.

Syntax Description	value	Route local-pref: 1-2147483648.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set local-preference 10`	
Related Commands	N/A	
Note		

set metric

set metric <value>
no set metric

Configures a route's entry metric parameter.
The no form of the command removes a route-map's entry metric setting.

Syntax Description	value	Route metric: 1-2147483647.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set metric 10`	
Related Commands	N/A	
Note		

set origin

set origin {egp | igp | incomplete}
no set origin

Configures a route's entry origin parameter.
The no form of the command removes a route-map's entry origin setting.

Syntax Description	egp	Set a route's entry origin parameter to external.
	igp	Set a route's entry origin parameter to internal.
	incomplete	Set a route's entry origin parameter to incomplete.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Role	admin	
Example	`switch (config route-map mymap permit 10)# set origin egp`	
Related Commands	N/A	
Note		

set tag

set tag <value>
no set tag

Configures a route's entry tag parameter.
The no form of the command removes a route-map's entry tag setting.

Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.4.0000	Updated parameter range
Role	admin	
Example	`switch (config route-map mymap permit 10)# set tag 10`	
Related Commands	N/A	
Note		

set weight

set weight <number>
no set weight

Configures modifications to redistributed routes.
The no form of the command removes this configuration from the route map clause.

Syntax Description	number	Value of the weight to set. Range: 1-65535.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
	3.4.0000	Updated parameter range
Role	admin	
Example	`switch (config route-map mymap permit 10)# set weight 7`	
Related Commands	route-map <map-name> [deny | permit] [sequence-number]	
Note		

show route-map

show route-map [<name>]

Displays route map configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.5006
Role	admin

Example	`switch (config)# show route-map mymap`
	`route-map mymap, permit, sequence 1200`
	`  Set clauses:`
	`    continue 1800`
	`switch (config)#`

Related Commands	N/A
Note	

6.6.2 IP Prefix-List

Prefix-list is a list of entries, each of which can match one or more IP prefixes. A prefix-list is usually used to match a specific IP prefix, mostly in relation to IP route destinations.

The prefix is considered to match the list if one of the entries match the prefix; the entry itself can be marked as a “permit” entry or a “deny” entry, which can be used by the matching code to decide if the route is to be accepted or not.

The prefix is matched to the prefix-list entries in the order of the sequence number of the entries in the list.

6.6.2.1 Commands

ip prefix-list

ipv6 prefix-list <list-name> [seq <number>] {permit | deny} <ipv6> <length> [eq <length> | le <length> | ge <length> [le <length>]]]
no ipv6 prefix-list <list-name> [seq <number>]

The command creates or updates IPv6 prefix-list.
The no form of the command deletes the prefix-list or a prefix-list entry

Syntax Description	list-name	String
	seq <number>	Sequence number assigned to entry. Range: 0-65535.
	permit	Permits access for a matching condition.
	deny	Denies access for a matching condition.
	ipv6	IPv6 address
	length	Prefix length
	eq | ge | le <mask>	• eq: Equal to a specified prefix length • ge: Greater than or equal to a specified prefix length • le: Less than or equal to a specified prefix length
Default	Sequence value = 10	
Configuration Mode	config	
History	3.3.5200	
	3.6.4070	Added support for IPv6
Role	admin	
Example	`switch (config)# ipv6 prefix-list a-list permit 2001::0 /64 eq 32` `switch (config)#`	
Related Commands	N/A	
Note		

show ipv6 prefix-list

show ipv6 prefix-list [<name>]

Displays IPv6 prefix-lists.

Syntax Description	name	Displays a specific prefix-list.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.4070	Added support for IPv6
Role	admin	

Example	`switch (config)# show ipv6 prefix-list`
	`prefix-list: a-list`
	`count: 1, range entries: 1, sequences: 10 - 10`
	`seq 10 permit 2001::0 /64 ge eq 32 (hit count: 0, refcount: 0)`
	
	`switch (config)#`	

Related Commands	N/A	
Note		

6.7 Multicast (IGMP and PIM)

Protocol independent multicast (PIM) is a collection of protocols that deal with efficient delivery of IP multicast (MC) data. Those protocols are published in the series of RFCs and define different ways and aspects of multicast data distribution. PIM protocol family includes PIM dense mode (PIM-DM), Bootstrap router (BSR) protocol, and PIM sparse mode (PIM-SM) and Bidirectional PIM (PIM-BIDIR)—both of which are not supported on Onyx.

PIM builds and maintains multicast routing tables based on the unicast routing information provided by unicast routing tables that can be maintained statically or dynamically by IP routing protocols like OSPF and BGP.

6.7.1 Basic PIM-SM

PIM relies on the underlying topology gathering protocols that collect unicast routing information and build multicast routing information base (MRIB). The primary role of MRIB is to determine the next hop for PIM messages. MC data flows along with the reverse path of the PIM control.

MC tree construction contains three phases:

1. Construction of a shared distribution tree. This tree is built around a special designated router (DR) called the rendezvous point (RP).
2. Establishing a native forwarding path from MC sources to the RP
3. Building an optimized MC distribution tree from each MC source to all MC targets bypassing the RP

The first stage of the multicast tree establishment starts when the MC receiver expresses desire to start receiving MC data. It can happen as a result of using one of the L2 protocols like MLD or IGMP, or by static configuration. When such request is received by the last hop router (a designated router) this router starts to build a distribution path from the RP. It starts to send periodic "Join" messages to the nearest PIM neighbor router towards the RP. The next router continues to do the same. Eventually the process converges when Join messages reach RP or a router that has already created that distribution tree. Usually that tree is called a shared tree because it is created for any source for specific MC group G and is noted as (*,G).

At that stage, MC senders can start sending MC data. The DR next to the MC source extracts the packets from the data flow and tunnels them to the RP. The RP decapsulates the packets and distributes them to all MC receivers along with the share tree.

On the second stage the RP switches from tunneling of multicast packets from MC sources to forwarding native traffic. When the RP identifies that a new MC source started to send packets, it initiates an establishment of a native forwarding path from the DR of that source to itself. For this purpose it starts to send Join messages towards MC source to nearest neighbor to that source according the MRIB. This is a source specific Join and is noted as (S,G). When data path is established up to the DR, the DR switches from tunneling MC packets to their native forwarding, so the RP does not need to decapsulate MC packets anymore, but still continue to distribute the packets along with shared tree.

On the third phase multicast receivers will try to switch from shared tree to source specific tree by creating a direct distribution path from a multicast source. When last hop router of the multicast receiver identifies multicast traffic coming from any multicast source it will start to send

Join messages towards the source with purpose to create a direct source specific path to that source. Once such path will be established and Designated router that is attached to the source L2 network will start to distribute the multicast traffic directly bypassing shared tree, the last hop router will detach its receivers from shared tree for that data and will switch to the shortest path tree distribution.

6.7.2 Source-Specific Multicast (SSM)

Source-Specific Multicast (SSM) is a method of delivering multicast packets in which the only packets that are delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security.

SSM requires that the receiver specify the source address and explicitly excludes the use of the (*,G) join for all multicast groups in RFC 3376, which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.

Source-specific multicast is best understood in contrast to any-source multicast (ASM). In the ASM service model a receiver expresses interest in traffic to a multicast address. The multicast network must discover all multicast sources sending to that address, and route data from all sources to all interested receivers.

This behavior is particularly well suited for groupware applications where all participants in the group want to be aware of all other participants, and the list of participants is not known in advance.

The source discovery burden on the network can become significant when the number of sources is large.

In the SSM service model, in addition to the receiver expressing interest in traffic to a multicast address, the receiver expresses interest in receiving traffic from only one specific source sending to that multicast address. This relieves the network of discovering many multicast sources and reduces the amount of multicast routing information that the network must maintain.

SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components, including routers and even the sending host. Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in RFC 4607.

SSM destination addresses must be in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6.

Source-specific multicast delivery semantics are provided for a datagram sent to an SSM address. That is, a datagram with source IP address S and SSM destination address G is delivered to each upper-layer “socket” that has specifically requested the reception of datagrams sent to address G by source S, and only to those sockets.

6.7.3 Bootstrap Router

For correct operation each PIM router requires a capability to map a multicast group that it needs to serve to a Rendezvous point for that group. This mapping can be done manually or the mapping can be distributed dynamically in the network. BSR protocol serves for this purpose.

This protocol introduces new role in the multicast network – Bootstrap router. That router is responsible to flood multicast group to RP mapping through the multicast routing domain. Boot-

strap router is elected dynamically among bootstrap router candidates (C-BSR) and once elected will collect from Rendezvous point candidate (C-RP) mapping information and distribute it in the domain.

Bootstrap activity contains 4 steps. First each C-BSR configured in the network originates floods into the network bootstrap messages that express the router desire to become BSR and also its BSR priority. Any C-BSR that receives that information and has lower priority will suspend itself, so eventually only one router will send BSR messages and become BSR.

When BSR is elected all RP candidates start to advertise to BSR a list of groups that this RP can serve. On the next step, after BSR learns the group mapping proposals, it forms a final group to RP mapping in the domain and starts to distribute it among PIM routers in the multicast routing domain. When PIM router receives BSR message with the group to RP mapping, it installs that mapping in the router local cache and uses that information to create multicast distribution trees.

6.7.4 Configuring Multicast

Precondition steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:

```
switch (config)# vlan 10
```

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)#  1/1
switch (config ethernet 1/1)#switchport access vlan 10
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

Step 5. Apply IP address to the VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.10 /24
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

6.7.4.1 Configuring IGMP

IGMP is enabled when IP multicast is enabled and static multicast or PIM is enabled on the interface.

6.7.4.2 Verifying IGMP

Step 1. Display a brief IGMP interface status. Run:

```
switch (config)# show ip igmp interface brief

IGMP Interfaces for VRF "default", count: 3

------------------------------------------------------------------------------
Interface     IP Address      IGMP Querier     Membership Count     Version
------------------------------------------------------------------------------
Vlan10        10.10.10.1      10.10.10.1       1                    v2
```

Step 2. Display detailed IGMP interface status. Run:

```
switch (config)# show ip igmp interface vlan 10
Interface vlan10
        Status: protocol-down/link-down/admin-up
        VRF: "vrf-default"
        IP address: 10.10.10.1/24
        Active querier: 10.10.10.1
        Version: 2
        Next query will be sent in: 00:01:45
        Membership count: 0
        IGMP version: 2
        IGMP query interval: 125 secs
        IGMP max response time: 10 secs
        IGMP startup query interval: 31 secs
        IGMP startup query count: 2
        IGMP last member query interval: 1 secs
        IGMP last member query count: 2
        IGMP group timeout: 260 secs
        IGMP querier timeout: 0 secs
        IGMP unsolicited report interval: 10 secs
        IGMP robustness variable: 2
        IGMP interface immediate leave: Disabled
        Multicast routing status on interface: Enabled
        Multicast TTL threshold: 0

        IGMP interface statistics:
          General (sent/received):
            v2-queries: 2/0
            v2-reports: 0/0
            v2-leaves : 0/0
            v3-queries: 0/0
            v3-reports: 0/0

          Errors:
            Checksum errors                       : 0
            Packet length errors                  : 0
            Packets with Local IP as source       : 0
            Source subnet check failures          : 0
            Query from non-querier                : 0
```

```
            Report version mismatch                  : 0
            Query version mismatch                   : 0
            Unknown IGMP message type                : 0
            Invalid v2 reports                       : 0
            Invalid v3 reports                       : 0
            Invalid leaves                           : 0
            Packets dropped due to router-alert check: 0
```

Step 3. Display the list of IGMP groups and their status. Run:

```
switch (config)# show ip igmp groups
IGMP Connected Group Membership
Type: S - Static, D - Dynamic

---------------------------------------------------------------------------------------------------------
Group Address       Type            Interface         Uptime           Expires          Last Reporter
---------------------------------------------------------------------------------------------------------
226.0.1.0           D               vlan10            00:00:05         N/A              10.10.10.2
226.0.1.1           D               vlan10            00:00:04         N/A              10.10.10.2
```

6.7.4.3 Configuring PIM

Prerequisites:

Step 1. If not enabled, enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. Globally enable multicast routing. Run:

```
switch (config)# ip multicast-routing
```

➢ ***To configure PIM:***

Step 1. Enable PIM. Run:

```
switch (config)# protocol pim
```

Step 2. Enable PIM on any IP interface (router port or VLAN interface) facing an L3 multicast source or L3 multicast receiver including transit interfaces. For example, run:

```
switch (config)#  1/4 ip pim sparse-mode
```

The interface's primary address is always used in PIM.

Step 3. Enable IGMP on any IP interface (router port or VLAN interface) facing multicast receivers. For example, run:

```
switch (config)#  1/4 ip igmp version {2|3}
```

If IGMP must be enabled on VLAN interface, IP IGMP snooping must also be enabled (globally and on the relevant VLAN):

```
switch (config)# interface vlan 50 ip igmp version {2|3}
switch (config)# ip igmp snooping
switch (config)# vlan 50 ip igmp snooping
```

Step 4. Configure a rendezvous point. Run:

```
switch (config)# ip pim rp-address 10.10.10.10
```

Generally, the RP address must be the same as the loopback address of the switch acting as the RP.

The RP address must be reachable to all switches.

Step 5. Configure a group mapping for a static RP. Run:

```
switch (config)# ip pim rp-address 192.168.0.1
```

You may also specify a “group-list <ip-address> <prefix>” parameter (ip pim rp-address 192.168.0.1 group-list 224.0.0.0/4) if you want different RPs for different groups.

6.7.5 Commands

6.7.5.1 PIM

protocol pim

	protocol pim **no protocol pim** Enables protocol independent multicast (PIM). The no form of the command hides all PIM commands and deletes all PIM configurations.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.5006
Role	admin
Example	`switch (config) # protocol pim`
Related Commands	N/A
Note	

ip pim sg-expiry-timer

ip pim sg-expiry-timer <seconds>
no ip pim sg-expiry-timer

Adjusts the SG expiry timer interval for PIM-SM SG multicast routes.
The no form of the command resets the parameters to their default values

Syntax Description	seconds	Range: 1-65535 seconds
Default	seconds: 180	
Configuration Mode	config	
History	3.6.6102	
Role	admin	
Example	`switch (config) # ip pim sg-expiry-timer 180`	
Related Commands	N/A	
Note		

ip pim rp-address

ip pim rp-address <rp-address> [group-list <ip-address> <prefix>] [override]
no ip pim rp-address <rp-address> [group-list <ip-address> <prefix>] [override]

Configures a static IP address of a rendezvous point for a multicast group range or adds new multicast range to existing RP.
The no form of the command removes the rendezvous point for a multicast group range or removes all configuration of the RP.

Syntax Description	rp-address	The static IP address of rendezvous point
	ip-address	IP address of the group-range (coupled with the prefix parameter)
	prefix	Network prefix (in the format of /24, or 255.255.255.0 for example) of group range
	override	Specifies that this configuration overrides dynamic configuration learned by BSR
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
Role	admin	
Example	`switch (config) # ip pim rp-address 10.10.10.10`	
Related Commands	N/A	
Note		

ip pim bsr-candidate

ip pim bsr-candidate {vlan <vlan-id> | loopback <number> | ethernet <port> | port-channel <id>} [hash-len <hash-length>] [priority <priority>] [interval <interval>]
no ip pim bsr-candidate {vlan <vlan-id> | loopback <number> | ethernet <port>} [hash-len <hash-length>] [priority <priority>] [interval <interval>]

Configures the switch as a candidate BSR router (C-BSR).
The no form of the command removes BSR-candidate configuration or restores default parameters values.

Syntax Description	vlan <vlan-id>	VLAN ID. Range is 1-4094.
	loopback <number>	Loopback interface for the BSR candidate address
	ethernet <port>	Ethernet interface for the BSR candidate address
	port-channel <id>	LAG interface for the BSR candidate address
	hash-len	Specifies the hash mask length used in BSR messages. Range: 0-32.
	priority	BSR priority rating. Larger numbers denote higher priority. Range: 0-255.
	interval	Period between the transmission of BSMs (seconds). Range:10-536870906.
Default	The interface is not BSR candidate by default. • priority: 64 • interval: 60 • hash-len: 30	
Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan	
History	3.3.5006	
Role	admin	
Example	`switch (config) # ip pim bsr-candidate vlan 10 priority 100`	

Related Commands	ip pim sparse-mode
Note	• IP PIM sparse-mode must be enabled on the interface. • A BSR is a PIM router within the PIM domain through which dynamic RP selection is implemented. The BSR selects RPs from a list of candidate RPs and exchanges bootstrap messages (BSM) with all routers in the domain. The BSR is elected from one of the C-BSRs through an exchange of BSMs. A subset of PIM routers within the domain are configured as candidate Bootstrap routers (C-BSRs). Through the exchange of Bootstrap messages (BSMs), the C-BSRs elect the BSR, which then uses BSMs to inform all domain routers of its status. • Command parameters specify the switch's BSR address, the interval between BSM transmissions, hash length used for RP calculations and the priority assigned to the switch when electing a BSR. • Entering an ip pim bsr-candidate command replaces any previously configured bsr-candidate command. If the new command does not specify a priority or interval, the previously configured values persist in running-config.

ip pim register-source

ip pim register-source <interface>
no ip pim register-source <interface>

Configures interface from which to use IP as source in PIM communications.
The no form of the command undoes this configuration.

Syntax Description	interface	Interface whose IP to use
Default	N/A	
Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan	
History	3.6.6102	
Role	admin	
Example	`switch (config) # ip pim register-source ethernet 1/2`	
Related Commands	N/A	
Note	This command must be set on an L3 interface with PIM sparse-mode (and not on a regular L3 interface which is not a PIM interface).	

ip pim rp-candidate

ip pim rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/port>} group-list <ip-address> <prefix> [priority <priority>] [interval <interval>]
no ip pim rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/port>} group-list <ip-address> <prefix> [priority <priority>] [interval <interval>]

Configures the switch as a candidate rendezvous point (C-RP).
The no form of the command removes the ip pim rp-candidate from running-config command for the specified multicast group.

Syntax Description	ethernet <slot/port>	Ethernet interface.
	port-channel <number>	LAG interface.
	vlan <vlan-id>	VLAN ID. Range: 1-4094.
	loopback <number>	Loopback interface number.
	ip-address	The group IP address.
	prefix	Network prefix (for example /24, or 255.255.255.0).
	priority	RP priority rating. Range: 0-255, where smaller numbers mean higher priority.
	interval	RP-advertisements message transmission interval. Range: 0-16383.
Default	The RP priority is 192. The BSR message interval is 60 seconds.	
Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan	
History	3.3.5006	
Role	admin	
Example	`switch (config) # ip pim rp-candidate vlan 19 group-list 225.6.5.0 /25 priority 20 interval 30`	

Related Commands	N/A
Note	• The BSR selects a multicast group's dynamic RP set from the list of C-RPs in the PIM domain. The command specifies the interface (used to derive the RP address), C-RP advertisement interval, and priority rating. The BSR selects the RP set by comparing C-RP priority ratings. The C-RP advertisement interval specifies the period between successive C-RP advertisement message transmissions to the BSR. • Running-config supports multiple multicast groups through multiple ip pim rp-candidate statements: • All commands must specify the same interface. Issuing a command with an interface that differs from existing commands removes all existing commands from running-config. • Running-config stores the interval and priority setting in a separate statement that applies to all rp-candidate statements. When a command specifies an interval that differs from the previously configured value, the new value replaces the old value and applies to all configured rp-candidate statements. The default interval value is 60 seconds. • When the no commands do not specify a multicast group, all rp-candidate statements are removed from running-config. The no ip pim rp-candidate interval commands restore the interval setting to the default value of 60 seconds. • When setting a priority, all previous rp-candidates within all interfaces and groups are configured to this priority.

ip pim sparse-mode

	ip pim sparse-mode **no ip pim sparse-mode** Sets PIM sparse mode on this interface. The no form of the command disables the sparse-mode on the interface and deletes all interfaces configuration.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.5006
Role	admin
Example	`switch (config interface vlan 10) # ip pim sparse-mode`
Related Commands	N/A
Note	

ip pim dr-priority

ip pim dr-priority <priority>
no ip pim dr-priority

Configures the designated router (DR) priority of PIM Hello messages.
The no form of the command resets this parameter to its default.

Syntax Description	priority	The designated router priority of the PIM Hello messages. Range is 1-4294967295.
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10) # ip pim dr-priority 5`	
Related Commands	ip pim sparse-mode	
Note	The command “ip pim sparse-mode” must be run prior to using this command.	

ip pim hello-interval

	ip pim hello-interval <interval> **no ip pim hello-interval** Configures PIM Hello interval in milliseconds. The no form of the command resets this parameter to its default.	
Syntax Description	interval	PIM Hello interval in seconds . Range:1-18000.
Default	30 seconds	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
	3.6.4006	Updated range
Role	admin	
Example	`switch (config interface vlan 10) # ip pim hello-interval 7000`	
Related Commands	ip pim sparse-mode	
Note	The command “ip pim sparse-mode” must be run prior to using this command.	

ip pim join-prune-interval

ip pim join-prune-interval <period>
no ip pim join-prune-interval

Configures the period between Join/Prune messages that the configuration mode interface originates and sends to the upstream RPF neighbor.
The no form of the command resets this parameter to its default.

Syntax Description	period	Range: 1-18000 seconds.
Default	60 seconds	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5200	
	3.6.4006	Updated range
Role	admin	
Example	`switch (config interface vlan 10) # ip pim join-prune-interval 60`	
Related Commands		
Note		

ip pim ssm range

ip pim ssm range {standard | group-list {<group-range>|<address> <prefix>}}
no ip pim ssm range {standard | group-list {<group-range>|<address> <prefix>}}

Enables one or more ranges for SSM operation.
The no form of the command disables range for SSM operation.

Syntax Description	standard	set the SSM operation to standard SSM range 232.0.0.0/8
	<group-range>	user-defined multicast range for SSM operation. Example 233.0.0.0/8
	<ip-address>	group range ip-address. Example: 233.0.0.0
	<prefix>	group range prefix. Example /8 or 255.0.0.0
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Role	admin	
Example	`switch (config) # ip pim ssm range group-list 234.0.0.0/8`	
Related Commands		
Note	Standard and group-list configurations are mutually exclusive: It is necessary to delete standard SSM configuration in order to add group-list and it is necessary to delete all existing group-list configuration in order to configure standard SSM configuration.	

ip pim multipath next-hop

ip pim [vrf <vrf-name>] multipath next-hop [<algorithm>]
no ip pim [vrf <vrf-name>] multipath next-hop

Configures PIM next-hop calculation algorithm.
The no form of the command resets PIM next-hops configuration to default (highest neighbor).

Syntax Description	vrf	VRF name
	algorithm	Selectable next-hop calculation algorithms: • g-hash – selects next-hop according to group address • mod – split groups between next hops on a module basis • s-g-hash - Selects next-hop according to group and source address
Default	Highest neighbor – next-hop with highest IP address is selected	
Configuration Mode	config	
History	3.6.8100	
	3.7.11xx	Updated syntax
Role	admin	
Example	`switch (config) # ip pim multipath next-hop g-hash`	
Related Commands	N/A	
Note		

ip pim multipath rp

ip pim multipath rp [<algorithm>]
no ip pim multipath rp

Configures PIM RP selection algorithm.
The no form of the command resets PIM RP selection algorithm to default (g-hash algorithm which is described in RFC 4601, sec. 4.7.2).

Syntax Description	algorithm	Selectable RP selection algorithms: • mod – split groups between RPs on a module basis
Default	g-hash - RPs are selected according to group address.	
Configuration Mode	config	
History	3.7.11xx	
Role	Admin	
Example	`switch (config) # ip pim multipath rp mod`	
Related Commands	N/A	
Note		

clear ip pim counters

	clear ip pim counters Clears PIM counter information
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6102
Role	admin
Example	`switch (config) # clear ip pim counters`
Related Commands	
Note	

show ip pim protocol

show ip pim [vrf {all | <vrf_name>}] protocol

Displays PIM protocol information:
1. Counters
2. Next-hop selection algorithm
3. RP selection algorithm
4. (S, G) expiry timer

Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated example output
	3.6.8008	Updated Example and added “vrf” parameter
	3.7.11xx	Updated description and example output
Role	admin	
Example		
Related Commands		
Note		

```
switch (config) # show ip pim vrf default protocol

PIM Control Counters for VRF "default":
Next-hop selection: highest neighbor
RP selection: hash4601
(S,G) expiry timer: 210 seconds

PIM Control Counters:
  ------------------------------------------------------------
  Counters             Received     Sent         Invalid
  ------------------------------------------------------------
  Assert               0            0            0
  Bootstrap Router     224          218          0
  CRP Advertisement    0            0            0
  Hello                443          551          0
  J/P                  0            0            0
  Join                 0            0            N/A
  Prune                0            0            N/A
  Register             0            0            0
  Register Stop        0            0            0
  State Refresh        0            0            0
```

show ip pim bsr

show ip pim [vrf {all | <vrf_name>}] bsr

Displays PIM BSR information.

Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example	`switch (config) # show ip pim vrf all bsr`
	
	`PIMv2 Bootstrap information for VRF "default":`
	`  No BSR is currently elected.`
	`  This system is not a candidate-BSR`
	
	`PIMv2 Bootstrap information for VRF "vrf_1":`
	`  BSR address          : 17.17.17.10`
	`  Uptime               : N/A`
	`  BSR Priority         : 64`
	`  Hash mask length     : 30`
	`  Expires              : 00:00:34`
	`  Candidate BSR        : Yes`
	`  Candidate BSR address: 17.17.17.10`
	`  priority             : 64`
	`  hash mask length     : 30`
	`  interval             : N/A`
	`  holdtime             : N/A`	

Related Commands		
Note		

show ip pim interface

show ip pim [vrf {all | <vrf_name>}] interface {[ethernet <port> | port-channel <id> | vlan <vlan id>]}

Displays information about the enabled interfaces for PIM.

Syntax Description	vrf	Displays output for a specific VRF
	ethernet <port>	Filters the output for specific Ethernet port
	port-channel <id>	Filters the output for specific LAG interface
	vlan <vlan-id>	Filters the output for specific VLAN interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated example output
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example

```
switch (config)# show ip pim vrf default  1/17

VRF "default":
  Interface eth1/17 address 17.17.17.10:
    PIM                       : enabled
    PIM version               : 2
    PIM mode                  : sparse
    PIM DR                    : 17.17.17.10 (this system)
    PIM DR Priority           : 1
    PIM configured DR priority: 1
    PIM neighbor count        : 1
    PIM neighbor holdtime     : 105 secs
    PIM Hello Interval        : 30 seconds, next hello will be sent in:
00:00:00
    PIM Hello Generation ID   : d674dec2
    PIM Join-Prune Interval   : 60 seconds
    PIM domain border         :

    PIM Interface Statistics:
      General (sent/received):
        Hellos     : 125 / 123
        JPs        : 7 / 164
        Asserts    : 0 / 0
        DF-Offers  : N/A / N/A
        DF-Winners:: N/A / N/A
        DF-Backoffs: N/A / N/A
        DF-Passes  : N/A / N/A

      Errors:
        Checksum errors                          : N/A
        Invalid packet types/DF subtypes         : N/A / 0
        Authentication failed                    : N/A
        Packets from non-neighbors               : 0
        JPs received on RPF-interface            : N/A
        (*,G) Joins received with no/wrong RP    : N/A / N/A
        (*,G)/(S,G) JPs received for Bidir groups: N/A
```

Related Commands

Note

show ip pim interface brief

show ip pim [vrf {all | <vrf_name>}] interface brief

Displays PIM information summary for all interfaces.

Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	
Example		
Related Commands		
Note		

```
switch (config)# show ip pim vrf all interface brief

VRF "default":
  ------------------------------------------------------------------------
  Address         Interface     Ver/      Nbr    Query  DR     DR
                                Mode      Count  Intvl  Prior
  ------------------------------------------------------------------------
  20.20.20.10     eth1/1        v2/S      0      30     1      20.20.20.10
  30.30.30.10     eth1/2        v2/S      0      30     1      30.30.30.10
  17.17.17.10     eth1/17       v2/S      1      30     1      17.17.17.10
```

show ip pim neighbor

show ip pim [vrf {all | <vrf_name>}] neighbor [vlan <vlan-id> | <other interfaces> | <ip-addr>]

Displays information about IPv4 PIM neighbors.

Syntax Description	vrf	Displays output for a specific VRF
	vlan <vlan-id>	Filters the output per specific VLAN ID.
	neighbor-addr	Filters the output per specific neighbor IP address.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	
Example	`switch (config) # show ip pim vrf default neighbor`	
	`VRF "default":`	
	`---------------------------------------------------------------------------`	
	`Neighbor Interface Uptime Expires Ver DR-Prio Mode BFD`	
	`---------------------------------------------------------------------------`	
	`17.17.17.5 eth1/17 01:08:07 00:01:38 v2 1 None`	
Related Commands		
Note		

show ip pim rp

show ip pim [vrf {all | <vrf_name>}] rp [<rp-address>]

Displays information about the rendezvous points (RPs) for PIM.

Syntax Description	vrf	Displays output for a specific VRF
	rp-address	Address of the rendezvous point
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	
Example		
Related Commands		
Note		

```
switch (config)# show ip pim vrf all rp

PIM RP Status Information for VRF "default":
  BSR: Not Operational

PIM RP Status Information for VRF "vrf_1":
  BSR        : 17.17.17.10
  expires    : 44
  priority   : 64
  hash-length: 30

  RP 17.17.17.10:
    expires  : 00:02:07
    RP-source: 17.17.17.10

    group ranges:
      225.0.0.0/24, priority: 192
```

show ip pim rp-hash

show ip pim rp-hash <group>

Displays an RP that is selected for the given group.

Syntax Description	group	A group address for RP calculation
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated version
	3.7.11xx	Updated example
Role	Admin	

Example	`switch (config) # show ip pim rp-hash 224.1.1.0`
	
	`VRF "default":`
	`RP 192.167.7.1, v2:`
	`  RP-source:`
	`  priority   : N/A`
	`  Uptime     : N/A`
	`  Expires    : N/A`	

Related Commands		
Note	RP is calculated according PIMv2 hash function as described in RFC 4601	

show ip pim rp-candidate

show ip pim [vrf {all | <vrf_name>}] rp-candidate

Displays information about RP candidate status.

Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6000	Updated Example
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example	`switch (config)# show ip pim vrf all rp-candidate`
	
	`VRF "default":`
	`  No RP candidates`
	
	`VRF "vrf_1":`
	`  RP 17.17.17.10:`
	`  Interface  : eth1/17`
	`  Interval  : 60`
	`  Next advertisement in: 6`
	`  Holdtime  : 150`
	`  Priority  : 192`
	
	`  Group prefixes:`
	`  1: 225.0.0.0/24`	

Related Commands		
Note		

show ip pim ssm range

show ip pim ssm [vrf {all | <vrf_name>}] range

Displays information about configured PIM SSM ranges.

Syntax Description	vrf	Displays information about configured PIM SSM ranges per specified VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example	`switch (config)# show ip pim vrf all ssm range`
	
	`VRF "default":`
	`  PIM SSM is not configured`
	
	`VRF "vrf_1":`
	`  Range type  : group-list`
	`  Total number of entries: 1`
	
	`  Group ranges:`
	`  1: 234.1.1.0/24`
	`  2: 234.1.2.0/24`
	`  3: 234.1.3.0/24`
	`  4: 234.1.4.0/24`
	`  5: 234.1.5.0/24`	

Related Commands		
Note		

show ip pim upstream joins

show ip pim [vrf {all | <vrf_name>}] upstream joins

Displays information about any PIM joins/prunes which are currently being sent to upstream PIM routers

Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example	`switch (config) # show ip pim vrf all upstream joins`
	
	`VRF "default":`
	`  There are no upstream joins`
	
	`VRF "vrf_1":`
	`  Neighbor address 17.17.17.5:`
	`  via interface  : 17.17.17.10`
	`  next message in: N/A seconds`
	
	`  Group 238.0.0.1:`
	`  Joins:`
	`  1: 10.10.10.5`
	
	`  Prunes:`
	`  No prunes included`
	
	`  Group 225.0.0.1:`
	`  Joins:`
	`  1: 10.10.10.5`
	
	`  Prunes:`
	`  No prunes included`	

Related Commands		
Note	Should contain the following information: neighbor address, interface address, group range, Joins, Prunes.	

6.7.5.2 Multicast

ip multicast-routing

ip multicast-routing [vrf <vrf-name>]
no ip multicast-routing [vrf <vrf-name>]

Allows the switch to forward multicast packets.
The no form of the command disables multicast routing.

Syntax Description	vrf	VRF name
Default	Disabled	
Configuration Mode	config	
History	3.3.5006	
Role	admin	
Example	`switch (config)# ip multicast-routing`	
Related Commands	N/A	
Note		

ip mroute

ip mroute {<ip-addr> <ip-mask> <next-hop>} [pref]
no ip mroute {<ip-addr> <ip-mask>} [<next-hop>]

Configure multicast reverse path forwarding (RPF) static routes.
The no form of the command deletes the static multicast route.

Syntax Description	ip-addr	Unicast IP address.
	ip-mask	Network mask in a dotted format (e.g. 255.255.255.0) or /24 format.
	next-hop	Next hop IP address.
	preference	Route preference. Range: 1-255.
Default	Preference is 1	
Configuration Mode	config	
History	3.3.5006	
	3.6.6000	Added “next-hop” parameter to “no” form
Role	admin	
Example	`switch (config) # no ip mroute 2.1.1.0 /24 3.1.1.1`	
Related Commands	N/A	
Note		

ip multicast ttl-threshold

ip multicast ttl-threshold <ttl-value>
no ip multicast ttl-threshold

Configures the time-to-live (TTL) threshold of packets being forwarded out of an interface.
The no form of the command removes RPF static routes.

Syntax Description	ttl-value Range: 0-225.
Default	0 – all packets are forwarded
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.5006
Role	admin
Example	`switch (config interface vlan 10)# ip multicast ttl-threshold 10`
Related Commands	N/A
Note	

clear ip mroute

clear ip mroute [vrf <vrf>] [<group-address> [<source-address>]]

Clears multicast route information

Syntax Description	vrf	Clears multicast route information for specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Role	admin	
Example	`switch (config) # clear ip mroute 237.0.0.1 1.1.1.8`	
Related Commands	N/A	
Note	This command does not support clearing specific (S,G) state if G belongs to an ASM group range. Here (S,G) refers to source and group parameters accordingly.	

show ip mroute

show ip mroute [vrf {all | <vrf-name>}] [<group> [<prefix> [<source>]]]

Displays information about IPv4 multicast routes.

Syntax Description	source	Source IP address
	group	IP address of multicast group
	prefix	Network prefix of multicast group (in the format of /24, or 255.255.255.0 for example)
	summary	Displays a summary of the multicast routes
	vrf	Displays information pertinent to specified or all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000	
	3.5.1000	Added new F flag and updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	

Example

```
switch (config) # show ip mroute vrf vrf_1

IP Multicast Routing Table:
  Flags:
    B: Bidir Group
    A: ASM Group
    S: SSM Group
    L: Local
    P: Pruned
    R: RP-bit set
    T: SPT-bit set
    J: Join SPT
    F: Failed to install in H/W

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:04:40, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:04:40/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:04:37/00D 00:00:22, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 225.0.0.2/32), 00D 00:04:31, flags: A:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 225.0.0.3/32), 00D 00:04:16, flags: A:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 238.0.0.1/32), 00D 00:04:40/00D 00:00:19, flags: ST:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
         eth1/2, N/A/SSM, 00D 00:04:40/00D 00:00:00
```

```
show ip mroute vrf vrf_1 225.0.0.1

IP Multicast Routing Table:
  Flags:
    B: Bidir Group
    A: ASM Group
    S: SSM Group
    L: Local
    P: Pruned
    R: RP-bit set
    T: SPT-bit set
    J: Join SPT
    F: Failed to install in H/W

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:13:27, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:13:27/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:13:24/00D 00:00:35, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:

show ip mroute vrf all 225.0.0.1 /32

IP Multicast Routing Table:
  Flags:
    B: Bidir Group
    A: ASM Group
    S: SSM Group
    L: Local
    P: Pruned
    R: RP-bit set
    T: SPT-bit set
    J: Join SPT
    F: Failed to install in H/W

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:14:54, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:14:54/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:14:51/00D 00:00:08, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
```

Related Commands	N/A
Note	

show ip mroute summary

show ip mroute [vrf {all | <vrf-name>}] summary

Displays a summary of the IPv4 multicast routes.

Syntax Description	vrf	Displays information pertinent to specified or all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip mroute vrf vrf_1 summary

IP Multicast Routing Table:
  Flags:
    B: Bidir Group
    A: ASM Group
    S: SSM Group
    L: Local
    P: Pruned
    R: RP-bit set
    T: SPT-bit set
    J: Join SPT
    F: Failed to install in H/W

  Timers         : Uptime/Expires
  Interface state: Interface, Next-Hop or VCD, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32):
    Uptime   : 00D 00:11:18
    RP       : 17.17.17.10
    OIF count: 1
    flags    : AL

  (10.10.10.5, 225.0.0.1/32):
    Uptime   : 00D 00:11:15
    Exptime  : 00D 00:00:44
    OIF count: 0
    flags    : AT

  (10.10.10.5, 238.0.0.1/32):
    Uptime   : 00D 00:11:18
    Exptime  : 00D 00:00:41
    OIF count: 1
    flags    : ST

  Total: 3 routes
```

Related Commands	N/A
Note	

6.7.5.3 IGMP

ip igmp immediate-leave

ip igmp immediate-leave
no ip igmp immediate-leave

Enables the device to remove the group entry from the multicast routing table immediately upon receiving a leave message for the group.
The no form of the command disables immediate-leave.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.6.8100
Role	admin
Example	`switch (config interface vlan 10)# ip igmp immediate-leave`
Related Commands	N/A
Note	

ip igmp last-member-query-response-time

ip igmp last-member-query-response-time <interval>
no ip igmp last-member-query-response-time

Configures the IGMP last member query response time in seconds.
The no form of the command resets this parameter to its default.

Syntax Description	interval	IGMP last member query response time. Range:1-25 seconds.
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
	3.7.11xx	Updates note
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp last-member-query-response-time 10`	
Related Commands	N/A	
Note	• When both 'IGMP and IGMP Snooping' protocols handle a Leave message and have different values for "Last Member Query Time" timer configured, then there is traffic loss for a short period of time.	

ip igmp startup-query-count

ip igmp startup-query-count <count>
no ip startup-query-count

Configures the number of query messages an interface sends during startup. The no form of the command resets this parameter to its default.

Syntax Description	count	Range: 1-255
Default	2	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp startup-query-count 10`	
Related Commands	N/A	
Note		

ip igmp startup-query-interval

ip igmp startup-query-interval <interval>
no ip startup-query-interval

Configures the IGMP startup query interval in seconds.
The no form of the command resets this parameter to its default.

Syntax Description	interval	Range: 1-1800 seconds.
Default	31	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp startup-query-interval 10`	
Related Commands	N/A	
Note		

ip igmp query-interval

ip igmp query-interval <interval>
no ip igmp query-interval

Configures the IGMP query interval in seconds.
The no form of the command resets this parameter to its default.

Syntax Description	interval	The IGMP query interval. Range: 1-1800 seconds.
Default	125	
Configuration Mode	config interface vlan	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp query-interval 60`	
Related Commands	N/A	
Note		

ip igmp query-max-response-time

ip igmp query-max-response-time <time>
no ip igmp query-max-response-time

Configures the IGMP max response time in seconds.
The no form of the command resets this parameter to its default.

Syntax Description	time	The IGMP max response time. Range: 1-25 seconds.
Default	10	
Configuration Mode	config interface vlan	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp query-max-response-time 20`	
Related Commands	N/A	
Note		

ip igmp robustness-variable

ip igmp robustness-variable <count>
no ip igmp robustness-variable

Configures the IGMP robustness variable.
The no form of the command resets this parameter to its default.

Syntax Description	count	IGMP robustness variable. Range: 1-7.
Default	2	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp robustness-variable 4`	
Related Commands	N/A	
Note	• The robustness variable can be increased to increase the number of times that packets are resent. • This parameter reflects expected packet loss on a congested network.	

ip igmp static-oif

ip igmp static-oif <group> [source-ip <address>]
no ip igmp static-oif <group> [source-ip <address>]

Statically binds an IP interface to a multicast group.
The no form of the command deletes the static multicast address from the interface.

Syntax Description	group	Multicast IP address
	source-ip	IP address from which to receive group traffic
Default	N/A	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Role	admin	
Example	`switch (config interface vlan 10)# ip igmp static-oif 10.10.10.5`	
Related Commands	N/A	
Note	PIM must be enabled in order to configure the route in the hardware.	

clear ip igmp groups

clear ip igmp groups {all | interface <if> | vrf <number> | <group-address> <mask>}

Clears IGMP group information.

Syntax Description	all	Clears all IGMP groups
	interface	Clears IGMP groups on specific interface
	vrf	Clears IGMP groups in specific VRF
	group-address	Clears a specific group range
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
Role	admin	
Example	`switch (config)# clear ip igmp groups all`	
Related Commands	N/A	
Note		

show ip igmp groups

show ip igmp [vrf {all |<vrf_name>}] groups [<group> | <iface>]

Displays information about IGMP-attached group membership.

Syntax Description	vrf	Displays output for a specific VRF
	group	Filters the output to a specific IP multicast group address
	iface	Filters the output to a specific IP interface (i.e. ethernet, port-channel, vlan interface)
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	
Example		

```
switch (config)# show ip igmp vrf all groups

IGMP Connected Group Membership
Type: S - Static, D - Dynamic

VRF "default":
  No IGMP group memberships learned or configured

VRF "vrf_1":
  ------------------------------------------------------------------------------------------
  Group Address        Type        Interface        Uptime        Expires        Last Reporter
  ------------------------------------------------------------------------------------------
  225.0.0.1            D           eth1/1           01:03:03      00:03:51       20.20.20.5
  238.0.0.1            D           eth1/2           01:03:03      N/A            30.30.30.5
```

Related Commands	N/A
Note	

show ip igmp interface

show ip igmp [vrf {all | <vrf_name>}] interface [ethernet <if> | port-channel <if> | vlan <vlan-id>]

Displays IGMP brief configuration and status.

Syntax Description	brief	Displays brief output information
	ethernet	Displays output for a specific Ethernet port
	port-channel	Displays output for a specific LAG
	vlan <vlan-id>	Displays output for a specific VLAN ID
	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
	3.6.8100	Added “IGMP interface immediate leave” line to output
Role	admin	

Example

```
switch (config)# show ip igmp interface vlan 10
Interface vlan10
        Status: protocol-down/link-down/admin-up
        VRF: "vrf-default"
        IP address: 10.10.10.1/24
        Active querier: 10.10.10.1
        Version: 2
        Next query will be sent in: 00:01:45
        Membership count: 0
        IGMP version: 2
        IGMP query interval: 125 secs
        IGMP max response time: 10 secs
        IGMP startup query interval: 31 secs
        IGMP startup query count: 2
        IGMP last member query interval: 1 secs
        IGMP last member query count: 2
        IGMP group timeout: 260 secs
        IGMP querier timeout: 0 secs
        IGMP unsolicited report interval: 10 secs
        IGMP robustness variable: 2
        IGMP interface immediate leave: Disabled
        Multicast routing status on interface: Enabled
        Multicast TTL threshold: 0

        IGMP interface statistics:
          General (sent/received):
            v2-queries: 2/0
            v2-reports: 0/0
            v2-leaves : 0/0
            v3-queries: 0/0
            v3-reports: 0/0

          Errors:
            Checksum errors                         : 0
            Packet length errors                    : 0
            Packets with Local IP as source         : 0
            Source subnet check failures            : 0
            Query from non-querier                  : 0
            Report version mismatch                 : 0
            Query version mismatch                  : 0
            Unknown IGMP message type               : 0
            Invalid v2 reports                      : 0
            Invalid v3 reports                      : 0
            Invalid leaves                          : 0
            Packets dropped due to router-alert check: 0
```

Related Commands N/A

Note

show ip igmp interface brief

show ip igmp interface [ethernet <if> | port-channel <if> | vlan <vlan-id>] brief

Displays brief IGMP configuration and status information.

Syntax Description	vrf	Displays output for a specific VRF
	ethernet	Displays output for a specific Ethernet port
	port-channel	Displays output for a specific LAG
	vlan <vlan-id>	Displays output for a specific VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated Example
	3.6.8008	Updated Example and added “vrf” parameter
Role	admin	
Example	`switch (config)# show ip igmp vrf all interface brief`	
Related Commands	N/A	
Note		

```
VRF "default":
  ------------------------------------------------------------------------------------
  Interface          IP Address          IGMP Querier        Membership Count     Version
  ------------------------------------------------------------------------------------
  eth1/10            12.14.192.5         0.0.0.0             0                    v3

VRF "vrf_1":
  ------------------------------------------------------------------------------------
  Interface          IP Address          IGMP Querier        Membership Count     Version
  ------------------------------------------------------------------------------------
  eth1/1             20.20.20.10         20.20.20.10         1                    v2
  eth1/2             30.30.30.10         30.30.30.10         1                    v3
  eth1/17            17.17.17.10         17.17.17.5          0                    v3
```

6.8 VRRP

The Virtual Router Redundancy Protocol (VRRP) is a computer networking protocol that provides for automatic assignment of available IP routers to participating hosts. This increases the availability and reliability of routing paths via automatic default gateway selections on an IP subnetwork.

The protocol achieves this by creating virtual routers, which are an abstract representation of multiple routers (that is, a master and backup routers, acting as a group). The default gateway of a participating host is assigned to the virtual router instead of a physical router. If the physical router that is routing packets on behalf of the virtual router fails, another physical router is selected to automatically replace it. The physical router that is forwarding packets at any given time is called the master router.

VRRP provides information on the state of a router, not the routes processed and exchanged by that router. Each VRRP instance is limited, in scope, to a single subnet. It does not advertise IP routes beyond that subnet or affect the routing table in any way.

Routers have a priority of between 1-255 and the router with the highest priority becomes the master. The configurable priority value ranges from 1-254, the router which owns the interface IP address as one of its associated IP addresses has the priority value 255. When a planned withdrawal of a master router is to take place, its priority can be lowered, which means a backup router will preempt the master router status rather than having to wait for the hold time to expire.

6.8.1 Load Balancing

To create load balancing between routers participating in the same VR, it is recommended to create 2 (or more) VRs. Each router will be a master in one of the VRs, and a backup to the other VR(s). A group of hosts should be configured with Router 1's virtual address as the default gateway, while the second group should be configured with Router 2's virtual address.

Figure 41: Common VRRP Configuration with Load Balancing

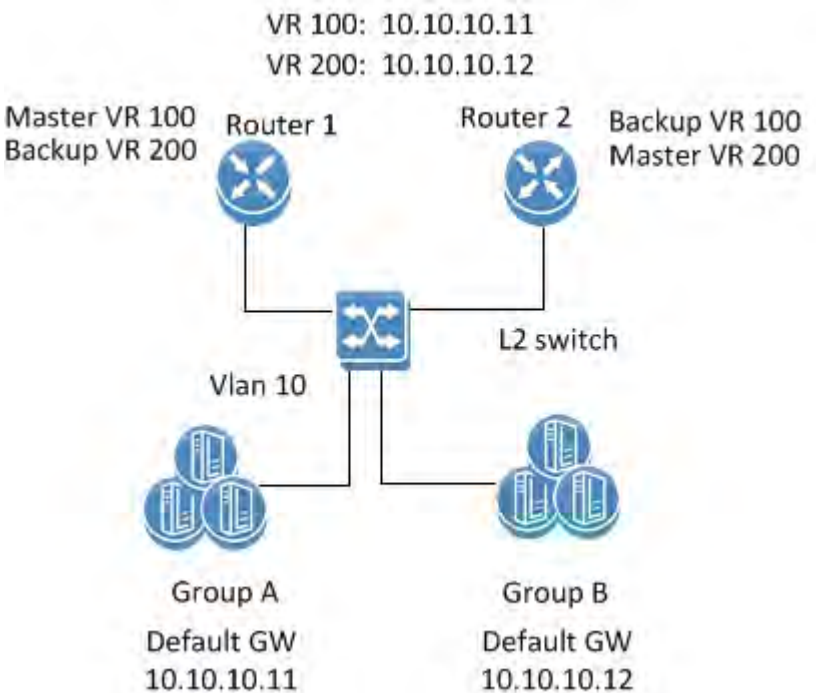

6.8.2 Configuring VRRP

➢ ***To configure VRRP:***

Precondition steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:

```
switch (config)# vlan 20
```

The VLAN cannot be the same one configured for the MLAG IPL, if MLAG is used.

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)#  1/1
switch (config ethernet 1/1)# switchport access vlan 20
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
```

Step 5. Apply IP address to the VLAN interface.

On one of the switches, run:

```
switch (config interface vlan 20)# ip address 20.20.20.20 /24
```

On the other switch, run:

```
switch (config interface vlan 20)# ip address 20.20.20.30 /24
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

Configure VRRP:

This is the same configuration on both switches

Step 1. Enable VRRP protocol globally. Run:

```
switch (config)# protocol vrrp
```

Step 2. Create a virtual router group for an IP interface. Up to 255 VRRP IDs are supported. Run:

```
switch (config interface vlan 20)# vrrp 100
```

Step 3. Set the VIP address. Run:

```
switch (config interface vlan 20 vrrp 100)# address 20.20.20.40
```

Step 4. Influence the election of the master in the VR cluster make sure that the priority of the desired master is the highest. Note that the higher IP address is selected in case the priority of the routers in the VR are the same. Select the priority. Run:

```
switch (config interface vlan 20 vrrp 100)# priority 200
```

Step 5. The advertizement interval should be the same for all the routers within the VR. Modify the interval. Run:

```
switch (config interface vlan 20 vrrp 100)# advertisement-interval 2
```

Step 6. The authentication text should be the same for all the routers within the VR. Configure the authentication text. Run:

```
switch (config interface vlan 20 vrrp 100)# authentication text my-password
```

Step 7. Use the `preempt` command to enable a high-priority backup virtual router to preempt the low-priority master virtual router. Run:

```
switch (config interface vlan 20 vrrp 100)# preempt
```

Step 8. Disable VRRP. Run:

```
switch (config interface vlan 20 vrrp 100)# shutdown
```

The configuration will not be deleted, only the VRRP state machine will be stopped.

6.8.3 Verifying VRRP

Step 1. Display VRRP brief status. Run:

```
switch (config)# show vrrp
Interface   VR  Pri  Time   Pre   State VR   IP addr
-----------------------------------------------------
Vlan20      1   200  2s     Y     Init       20.20.20.20
…
switch(config)#
```

Step 2. Display VRRP detailed status. Run:

```
switch (config)# show vrrp detail

VRRP Admin State : Enabled

Vlan20 - Group 1 (IPV4)

Instance Admin State : Enabled
State : Backup
Virtual IP Address : 20.20.20.40
Priority : 200
Advertisement interval (sec) : 2
Preemption : Enabled
Virtual MAC address : AA:BB:CC:DD:EE:FF

switch (config)#
```

Step 3. Display VRRP statistic counters. Run:

```
switch (config)# show vrrp statistics
Ethernet1/5 - Group 1 (IPV4)
Invalid packets:              0
Too short:                    0
Transitions to Master         6
Total received:               155
Bad TTL:                      0
Failed authentication:        0
Unknown authentication:       0
Conflicting authentication:   0
Conflicting Advertise time:   0
Conflicting Addresses:        0
Received with zero priority:  3
Sent with zero priority:      3
```

6.8.4 Commands

protocol vrrp

	protocol vrrp **no protocol vrrp** Enables VRRP globally and unhides VRRP related commands. The no form of the command deletes all the VRRP configuration and hides VRRP related commands.
Syntax Description	N/A
Default	no feature vrrp
Configuration Mode	config
History	3.3.4500
Role	admin
Example	`switch (config)# protocol vrrp`
Related Commands	
Note	

vrrp

vrrp <number>
no vrrp <number>

Creates a virtual router group on this interface and enters a new configuration mode. The no form of the command deletes the VRRP instance and the related configuration.

Syntax Description	number	A VRRP instance number. Range is 1-255.
Default	N/A	
Configuration Mode	config interface vlan	
History	3.3.4500	
	3.6.8100	Updated parameter range
	3.7.11xx	Updated Syntax & notes
Role	admin	
Example	`switch (config interface vlan 10)#` `switch (config interface vlan 10 vrrp 10)#`	
Related Commands		
Note	A maximum total of 64 VRRP instances are supported per switch system.	

address

address <ip-address> [secondary]
no address [<ip-address> [secondary]]

Sets virtual router IP address (primary and secondary).
The no form of the command deletes the IP address from the VRRP interface.

Syntax Description	ip-address	The virtual IP address.
	secondary	A secondary IP address for the virtual router.
Default	N/A	
Configuration Mode	config vrrp interface	
History	3.3.4500	
Role	admin	

Example	`switch (config vrrp 100)# address 10.10.10.10`
	`switch (config vrrp 100)# address 10.10.10.11 secondary`
	`switch (config vrrp 100)# address 10.10.10.12 secondary`	

Related Commands		
Note	• The virtual address can be either from the interface's primary or secondary subnet • This command is the enabler of the protocol. Therefore, set all the protocol parameters initially and only then set the ip-address. • There are up to 20 IP addresses associated with the VRRP instance. One primary and up to 19 secondary ip-addresses. • If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address (priority 255).	

shutdown

	shutdown **no shutdown** Disables the virtual router. The no form of the command enables the virtual router (stops the VRRP state machine).
Syntax Description	N/A
Default	Enabled (no shutdown)
Configuration Mode	config vrrp interface
History	3.3.4500
Role	admin
Example	`switch (config vrrp 100)# shutdown`
Related Commands	
Note	

priority

priority <level>
no priority

Sets the priority of the virtual router.
The no form of the command resets the priority to its default.

Syntax Description	level	The virtual router priority level. Range is 1-254.
Default	100	
Configuration Mode	config vrrp interface	
History	3.3.4500	
Role	admin	
Example	`switch (config vrrp 100)# priority 200`	
Related Commands		
Note	• The higher IP address will be selected as master, in case the priority of the routers in the VR are the same. • To influence the election of the master in the VR cluster make sure that the priority of the desired master is the higher.	

preempt

	preempt **no preempt** Sets virtual router preemption mode. The no form of the command disables the virtual router preemption.
Syntax Description	N/A
Default	Enabled (preempt)
Configuration Mode	config vrrp interface
History	3.3.4500
Role	admin
Example	`switch (config vrrp 100)# preempt`
Related Commands	
Note	To set this router as backup for the current virtual router master, preempt must be enabled.

authentication text

authentication text <password>
no authentication text

Sets virtual router authentication password and enables authentication.
The no form of the command disables the authentication mechanism.

Syntax Description	password	The virtual router authentication password. The password string must be up to 8 alphanumeric characters.
Default	Disabled	
Configuration Mode	config vrrp interface	
History	3.3.4500	
Role	admin	
Example	`switch (config vrrp 100)# authentication text mypassword`	
Related Commands		
Note		

advertisement-interval

advertisement-interval <seconds>
no advertisement-interval

Sets the virtual router advertisement-interval.
The no form of the command resets the parameter to its default.

Syntax Description	seconds	The virtual router advertisement-interval in seconds. Range: 1-255.
Default	1	
Configuration Mode	config vrrp interface	
History	3.3.4500	
Role	admin	
Example	`switch (config vrrp 100)# advertisement-interval 10`	
Related Commands		
Note		

show vrrp

show vrrp [interface <type> <number>] [vr <id>]

Displays VRRP brief configuration and status.

Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number.
	vr <id>	Filters the output to a specific virtual router. Range: 1-10.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
Role	admin	

Example	`switch(config)# show vrrp`
	`Interface  VR Pri Time  Pre  State VR  IP addr`
	`---------------------------------------------------`
	`Eth1/5     1  200  2s   Y    Init      192.0.1.10`
	`…`
	`switch(config)#`	

Related Commands		
Note		

show vrrp detail

show vrrp detail [interface <type> <number>] [vr <id>]

Displays detailed VRRP configuration and status.

Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number
	vr <id>	Filters the output to a specific virtual router. Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated Example
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config)# show vrrp detail`
	`VRRP Admin State: Enabled`
	
	`Vlan3200 - Vrrp 110 (IPV4):`
	`  Instance Admin State: Enabled`
	`  State: Init`
	`  Primary IP Address: 33.0.0.1`
	`  Virtual IP Address: 33.0.0.2`
	`  Priority: 100`
	`  Advertisement interval(sec): 2`
	`  Preemption: Enabled`
	`  Virtual MAC address: 00:00:5E:00:01:6E`
	`  Master router: 33.0.0.1`
	`  Master priority: 100`
	`  Master advertisement interval: 2`
	`  Associated IP Addresses:`
	`  33.0.0.3`
	`  33.0.0.4`	

Related Commands		
Note		

show vrrp statistics

show vrrp statistics [interface <type <number>] [vr <id>] [all]

Displays VRRP counters.

Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number
	vr <id>	Filters the output to a specific virtual router Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated Example
Role	admin	

Example	`switch (config)# show vrrp statistics`
	`VRRP Instance 100:`
	`  Invalid packets: 0`
	`  Too short: 0`
	`  Transitions to Master: 0`
	`  Total received: 0`
	`  Bad TTL: 0`
	`  Failed authentication: 0`
	`  Unknown authentication: 0`
	`  Conflicting authentication: 0`
	`  Conflicting Advertise time: 0`
	`  Conflicting Addresses: 0`
	`  Received with zero priority: 0`
	`  Sent with zero priority: 0`	

Related Commands		
Note		

6.9 MAGP

Multi-active gateway protocol (MAGP) is aimed to solve the default gateway problem when a host is connected to a set of switch routers (SRs) via MLAG.

The network functionality in that case requires that each SR is an active default gateway router to the host, thus reducing hops between the SRs and directly forwarding IP traffic to the L3 cloud regardless which SR traffic comes through.

6.9.1 Configuring MAGP

Prerequisite steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:

```
switch (config)# vlan 20
switch (config vlan 20)#
```

The VLAN cannot be the same one configured for the MLAG IPL, if MLAG is used.

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)#  1/1
switch (config  1/1)# switchport access vlan 20
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
switch (config interface vlan 20)#
```

Step 5. Set an IP address to the VLAN interface. Run:

```
switch (config interface vlan 20)# ip address 11.11.11.11 /8
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

➢ ***To configure MAGP:***

Step 1. Enable MAGP protocol globally. Run:

```
switch (config)# protocol magp
```

Step 2. Create a virtual router group for an IP interface. Run:

```
switch (config interface vlan 20)# magp 100
```

Up to 255 MAGP IDs are supported.

Step 3. Set a virtual router primary IP address. Run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router address 11.11.11.254
```

Only a virtual IP from the primary subnet can be configured for MAGP.

Step 4. Set a virtual router primary MAC address. Run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router mac-address
AA:BB:CC:DD:EE:FF
```

To obtain the virtual router's MAC address, please run the command "show vrrp detail".

➢ ***To verify the MAGP configuration, run:***

```
switch (config)# show magp 100
MAGP 100
  Interface vlan: 20
  Admin state: Master
  State: Enabled
  Virtual IP: 11.11.11.254
  Virtual MAC: AA:BB:CC:DD:EE:FF
```

This output is to be expected in both MAGP switches.

For more advanced configuration options, please refer to the following Mellanox Community post: https://community.mellanox.com/docs/DOC-1476.

6.9.2 Commands

protocol magp

	protocol magp **no protocol magp** Enables MAGP globally and unhides MAGP commands. The no form of the command deletes all the MAGP configuration and hides MAGP commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.4500
Role	admin
Example	`switch (config)# protocol magp` `switch (config)#`
Related Commands	
Note	IP routing must be enabled to enable MAGP.

magp

magp <instance>
no magp <instance>

Creates an MAGP instance on this interface and enters a new configuration mode. The no form of the command deletes the MAGP instance.

Syntax Description	instance	MAGP instance number. Range: 1-255.
Default	Disabled	
Configuration Mode	config interface vlan	
History	3.3.4500	
	3.7.11xx	Updated notes
Role	admin	
Example	`switch (config interface vlan 20)# magp 100` `switch (config interface vlan 20 magp 100)#`	
Related Commands		
Note	• Only one MAGP instance can be created on an interface • Different interfaces cannot share an MAGP instance • MAGP and VRRP are mutually exclusive • A maximum total of 64 MAGP instances are supported per switch system	

shutdown

	shutdown **no shutdown** Enables MAGP instance. The no form of the command disables the MAGP instance.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan magp
History	3.3.4500
Role	admin
Example	`switch (config interface vlan 10 magp 1)# shutdown`
Related Commands	
Note	

ip virtual-router address

ip virtual-router address <ip-address> [secondary]
no ip virtual-router address <ip-address> [secondary]

Sets MAGP virtual IP address.
The no form of the command resets this parameter to its default.

Syntax Description	ip-address	The virtual router IP address
	secondary	Adds secondary virtual router address
Default	N/A	
Configuration Mode	config interface vlan magp	
History	3.3.4500	
	3.6.8100	Added “secondary” parameter
Role	admin	
Example	`switch (config interface vlan 10 magp 1)# ip virtual-router address 10.10.10.10` `switch (config interface vlan 10 magp 1)#`	
Related Commands		
Note	The MAGP virtual IP address must be different from the interface IP address	

ip virtual-router mac-address

	ip virtual-router mac-address <mac-address> **no ip virtual-router mac-address** Sets MAGP virtual MAC address. The no form of the command resets the MAC address to its default.
Syntax Description	mac-address MAC address. Format: AA:BB:CC:DD:EE:FF.
Default	00:00:5E:00:01-<magp instance>
Configuration Mode	config interface vlan magp
History	3.3.4500
Role	admin
Example	`switch (config interface vlan 10 magp 1)# ip virtual-router mac-address AA:BB:CC:DD:EE:FF` `switch (config interface vlan 10 magp 1)#`
Related Commands	
Note	

show magp

show magp [<instance>]

Displays the MAGP configuration.

Syntax Description	instance	Displays configuration of a specific MAGP instance Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated Example
	3.6.8100	Updated Example
Role	admin	

Example	`switch (config)# show magp`
	
	`MAGP 1:`
	`  Interface vlan: 10`
	`  Admin state   : Enabled`
	`  State         : Master`
	`  Virtual IP    : 192.168.11.10`
	`  Virtual MAC   : 00:00:5E:00:01:14`
	
	`  Associated IP Addresses:`
	`  192.168.10.10`	

Related Commands		
Note		

show magp interface vlan

show magp interface vlan <id>

Displays the configuration of a specific MAGP instance.

Syntax Description	instance	MAGP instance number. Range: 1-255.
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated Example
	3.6.8100	Updated Example
Role	admin	

Example	`switch (config)# show magp interface vlan 10`
	
	`MAGP 1:`
	`  Interface vlan: 10`
	`  Admin state   : Enabled`
	`  State         : Master`
	`  Virtual IP    : 192.168.11.10`
	`  Virtual MAC   : 00:00:5E:00:01:14`
	
	`  Associated IP Addresses:`
	`  192.168.10.10`	

Related Commands		
Note		

6.10 DHCP Relay

Since Dynamic Host Configuration Protocol must work correctly even before DHCP clients have been configured, the DHCP server and DHCP client need to be connected to the same network.

In larger networks, this is not always practical because each network link contains one or more DHCP relay (DHCP-R) agents. These agents receive messages from DHCP clients and forward them to DHCP servers thus extending the reach of the DHCP beyond the local network.

DHCP-R is supported for IPv4 and IPv6.

DHCP-R is supported for both primary IP subnet and secondary IP subnets.

6.10.1 DHCP-R Virtual Routing and Forwarding (VRF) Auto-Helper

In some cases it is desired that DHCP-R functionality is automatically enabled to all IP interfaces in the system. For this purpose a vrf-auto-helper may be configured on a DHCP-R instance which would provide DHCP-R services automatically for each newly created interface on a VRF.

Only one instance in each VRF can have vrf-auto-helper capability. Whenever a new instance is created in a VRF, it automatically becomes a vrf-auto-helper.

It is possible to manually disable auto-helper capability for the instance. See command "vrf-auto-helper" on page 1589 for more information.

6.10.2 Upstream and Downstream Interfaces

It is possible to define an interface to be downstream, upstream, or bidirectional (both downstream and upstream):

- Bidirectional interface – capable of performing downstream and upstream functionalities
- Downstream interface (default configuration) – the interface on which queries are received from clients or from other relay agents
- Upstream interface – the interface to which queries from clients and other relay agents are forwarded

6.10.3 Commands

ip dhcp relay

ip dhcp relay [instance <instance-id>]
no ip dhcp relay [instance <instance-id>]

Enters DHCP relay instance configuration mode, and creates DHCP instance in active VRF context.
The no form of the command deletes the instance and DHCP relay process corresponding to it.

Syntax Description	instance-id Range: 1-8
Default	N/A
Configuration Mode	config
History	3.6.3004
Role	admin
Example	`switch (config)# ip dhcp relay instance 1` `switch (config ip dhcp relay instance 1)#`
Related Commands	N/A
Note	If an instance is not specified then instance 1 is used (if nonexistent, then it is created).

address

address <ip-address>
no address <ip-address>

Configures the DHCP server IP address on a particular instance.
The no form of the command deletes the DHCP server IP address.

Syntax Description	ip-address	Valid IP unicast address of DHCP server.
Default	N/A	
Configuration Mode	config ip dhcp relay	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
Role	admin	
Example	`switch (config ip dhcp relay instance 1)# address 1.2.3.4`	
Related Commands	ip dhcp relay	
Note	• Up to 16 IP addresses may be configured • To enable DHCP relay instance, at least one IP address should be configured, or always-on parameter should be turned on using the command “ip dhcp relay always-on” • The following option for running this command is also possible: ip dhcp relay instance 1 address <ip-address>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

always-on

always-on
no always-on

Enables broadcast mode on a particular instance.
The no form of the command disables the broadcast mode from instance.

Syntax Description	vrf	VRF name
Default	Disabled	
Configuration Mode	config ip dhcp relay	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
Role	admin	
Example	`switch (config ip dhcp relay instance 1)# always-on`	
Related Commands	ip dhcp relay	
Note	• Broadcasts DHCP requests to all interfaces with the DHCP relay agent for given VRF • In order to enable DHCP relay, at least one IP address should be configured, or always-on parameter should be turned on using this command • When DHCP servers are configured, requests are forwarded only to configured servers • The following option for running this command is also possible: ip dhcp relay instance 1 always-on. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

information option

	information option **no information option** Enables DHCP relay agents to insert option 82 on the packets of a particular instance. The no form of the command removes option 82 from the packets.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config ip dhcp relay
History	3.3.4150 3.6.3004 Enhanced command for DHCP-R multi-instance
Role	admin
Example	`switch (config ip dhcp relay instance 1)# information option`
Related Commands	ip dhcp relay
Note	The following option for running this command is also possible: ip dhcp relay instance 1 information option. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).

vrf

	vrf <vrf-name> **no vrf <vrf-name>** Configures mention instance in the given VRF. The no form of the command moves the instance back to default VRF.
Syntax Description	N/A
Default	N/A
Configuration Mode	config ip dhcp relay
History	3.6.3004
Role	admin
Example	`switch (config ip dhcp relay instance 1)# vrf 2`
Related Commands	N/A
Note	• If no VRF is specified, then the DHCP-R instance is created in the active VRF • If the VRF is changed, then the configuration of the DHCP-R instance is automatically deleted • The following option for running this command is also possible: ip dhcp relay instance 1 vrf <vrf-name>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).

port

port <udp-port>
no port <udp-port>

Changes the UDP port for the given instance.
The no form of the command sets the UDP port to default value.

Syntax Description	udp-port	UDP port Range: 1-65534
Default	67	
Configuration Mode	config ip dhcp relay	
History	3.6.3004	
Role	admin	
Example	`switch (config ip dhcp relay instance 1)# port 65534`	
Related Commands	N/A	
Note	• The system allocated 2 ports: One is the server port (udp-port), and another is client port (udp-port+1) • The following option for running this command is also possible: ip dhcp relay instance 1 port <udp-port>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

use-secondary-ip

	use-secondary-ip **no use-secondary-ip** Enables the switch to relay a single request from the client multiple times simultaneously, with each of the IP addresses configured on the corresponding downstream interfaces as the respective gateway address (linkaddr field of IPv4 DHCP request packet). The no form of the command disables this function.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config ip dhcp relay
History	3.6.8008
Role	admin
Example	`switch (config ip dhcp relay instance 1)# use-secondary-ip`
Related Commands	N/A
Note	

vrf-auto-helper

	vrf-auto-helper **no vrf-auto-helper** Makes all L3 interfaces (existing/newly created) to be part of the given instance. The no form of the command resets this parameter to its default
Syntax Description	N/A
Default	N/A
Configuration Mode	config ip dhcp relay
History	3.6.3004
Role	admin
Example	`switch (config ip dhcp relay instance 1)# vrf-auto-helper`
Related Commands	N/A
Note	• Every new DHCP-R instance created in a VRF automatically becomes the VRF auto-helper if no other DHCP-R instance has been configured VRF auto-helper previously in that VRF • The following option for running this command is also possible: ip dhcp relay instance 1 vrf-auto-helper. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).

ip dhcp relay instance (interface config)

ip dhcp relay instance <instance-id> [downstream] [upstream]
no ip dhcp relay instance <instance-id> [downstream] [upstream]

Enables the given interface to listen for DHCP packets coming from specified instance (i.e. binds interface to that instance).
The no form of the command removes the interface mapping from that instance.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	downstream	The interface on which queries are received from clients or from other relay agents
	upstream	The interface to which queries from clients and other relay agents should be forwarded
Default	Downstream	
Configuration Mode	config interface ethernet set as router port interface config interface port-channel config interface vlan	
History	3.6.3004	
	3.6.6000	Added downstream and upstream parameters
Role	admin	
Example	`switch (config  1/13)# ip dhcp relay instance 7 downstream`	
Related Commands	N/A	
Note	• In order to enable DHCP relay, other than configuring the downstream interface, at least one IP address must be configured, or the always-on parameter must be activated using the command “ip dhcp relay always-on” • When DHCP servers are configured, requests are forwarded only to configured servers • At most, 64 interfaces can be configured on each instance • Only an existent DHCP-R may be specified • Each interface is either upstream, downstream, or bidirectional • If only downstream interfaces are defined, all interfaces in VRF are assumed to be upstream interfaces	

clear ip dhcp relay counters

clear ip dhcp relay counters [vrf {<vrf-name> | all} | instance <instance-id>]

Clears all DHCP relay counters (all interfaces) in a given VRF or instance.

Syntax Description	vrf-name	VRF name or “all” for all VRFs
	instance-id	DHCP instance ID Range: 1-8
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
	3.6.5000	Added “all” parameter
Role	admin	
Example	`switch (config)# clear ip dhcp relay counters`	
Related Commands	N/A	
Note	• If no DHCP-R instance is specified, then the counters of all DHCP-R instances are cleared • If a VRF is specified, then the counters of all instances on that VRF are cleared • The command “clear counters all” may also be used to clear all DHCP-R counters	

6.10.3.1 Interface Commands

ip dhcp relay information option circuit-id

ip dhcp relay information option circuit-id <label>
no ip dhcp relay information option circuit-id

Specifies the content of the circuit ID sub-option attached to the client DHCP packet when it is forwarded a DHCP server.
The no form of the command removes the label assigned.

Syntax Description	label	Specifies the label attached to packets. The string may be up to 15 characters.
Default	The label is taken from the IP interface name (e.g. “vlan1”)	
Configuration Mode	config interface vlan config interface ethernet set as router port interface config interface port-channel set as router port interface	
History	3.3.4150	
	3.6.1002	Added VRF parameter
Role	admin	
Example	`switch (config interface vlan 10)# ip dhcp relay information options circuit-id my-label`	
Related Commands	N/A	
Note	The circuit ID sub-option is an IP interface attribute which is shared across all DHCP-R instances.	

6.10.3.2 Show Commands

show ip dhcp relay

show ip dhcp relay [instance <instance-id>]

Displays general DHCP configuration.

Syntax Description	instance-id	If instance ID is specified, then a particular instance configuration is displayed
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.1002	Added VRF and all parameters
	3.6.3004	Updated Example and parameters
	3.6.6000	Updated Example
	3.6.8008	Updated Example
Role	admin	

Example	`switch (config)# show ip dhcp relay`
	
	`Instance ID 1:`
	`  VRF Name: default`
	
	`  DHCP Servers:`
	`    1.1.1.1`
	
	`  DHCP relay agent options:`
	`    always-on         : Disabled`
	`    Information Option: Disabled`
	`    UDP port          : 67`
	`    Auto-helper       : Disabled`
	
	`  ----------------------------------------`
	`  Interface   Label            Mode`
	`  ----------------------------------------`
	`  eth1/5      N/A              downstream`	

Related Commands	N/A	
Note	• If no DHCP-R instance is given, then all DHCP-R instances are displayed • Only configured interfaces are displayed • Once vrf-auto-helper is enabled, no interface is displayed	

show ip dhcp relay counters

show ip dhcp relay counters [instance <instance-id> | vrf <vrf-name>]

Displays the DHCP relay counters.

Syntax Description	instance-id	Displays the DHCP relay counters for a given instance
	vrf	Displays the DHCP relay counters in a given VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.1002	Added VRF and all parameters
	3.6.5000	Updated Example
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ip dhcp relay counters

Instance 1:
  VRF Name: vrf-default

  DHCP Counter flags:
    SPR : Server Packets Received
    SPE : Server Packets Error
    SPRE: Server Packet Relayed
    CPR : Client Packets Received
    RP  : Relay Packets
    RE  : Relay Errors

  ---------------------------------
  Req/Resp    Received    Forwarded
  ---------------------------------
  All Req     0           0
  All Res     0           0

  -----------------------------------------------------
  If         SPRE       SPE        SPR        CPR
  -----------------------------------------------------
  eth1/5     0          0          0          0

  Packets Relayed to Server:
    ---------------------------------
    Server      RP          RE
    ---------------------------------
    1.1.1.1     0           0
```

Related Commands	N/A
Note	

6.10.4 DHCPv6 Relay

6.10.4.1 Commands

6.10.4.1.1Config

ipv6 dhcp relay instance

ipv6 dhcp relay instance <instance-id> [vrf-auto-helper] [downstream] [upstream]
no ipv6 dhcp relay instance <instance-id> [vrf-auto-helper]

Enables DHCP relay instance configuration mode, and creates DHCP instance in active VRF context.
An instance without an assigned addresses is sent to All_DHCP_servers address.
The no form of the command deletes the DHCP relay instance.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	vrf-auto-helper	Instance becomes VTF auto helper
	downstream	The interface on which queries are received from clients or from other relay agents
	upstream	The interface to which queries from clients and other relay agents should be forwarded
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.6.4070	First release
	3.6.6000	Added downstream and upstream parameters
Role	admin	
Example	`switch (config  1/1) # ipv6 dhcp relay instance 1 downstream`	
Related Commands		

Note	• Each interface is either upstream, downstream, or bidirectional
	• At most, 64 interfaces can be configured on each instance
	• If only downstream interfaces are defined, all interfaces in VRF are assumed to be upstream interfaces
	• An instance must meet two conditions to become active
	  • a server address or an upstream interface
	  • a downstream interface	

ipv6 dhcp relay instance (global server)

ipv6 dhcp relay instance <instance-id> address <ipv6-address or list of addresses>
no ipv6 dhcp relay instance <instance-id> address <ipv6-address or list of addresses>

Configure the server address on a particular instance.
Instance without assigned addresses will send to All_DHCP_servers address.
The no form of the command will delete the server address from instance.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	ipv6-address	Valid global unicast IPv6 server address Up to 16 addresses can be assigned per instance
Default	N/A	
Configuration Mode	config	
History	3.6.4070	First release
Role	admin	
Example	`switch (config)# ipv6 dhcp relay instance 1 address 2001::1`	
Related Commands		
Note		

ipv6 dhcp relay instance address (destination address on interface)

ipv6 dhcp relay instance <instance-id> address <link-local-address>
no ipv6 dhcp relay instance <instance-id> address <link-local-address>

Configures the destination address on a particular instance on a specific upstream interface. Only link local address is supported.
The no form of the command deletes the destination address on a specific upstream interface from a particular instance.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	ipv6-address	Destination unicast or multicast address Only link local address in supported
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.6.4070	First release
Role	admin	
Example	`switch (config  1/13)# ipv6 dhcp relay instance 1 address fe80::1`	
Related Commands		
Note	Up to 16 addresses can be assigned per instance	

ipv6 dhcp relay instance interface-id option

[no] ipv6 dhcp relay instance <instance-id> interface-id option

Enables the instance to insert interface ID option.
The no form of the command disables this option.

Syntax Description	instance-id	DHCP instance ID. Range: 1-8.
Default	Default interface-id is an interface name (e.g. vlan1, eth1/1)	
Configuration Mode	config	
History	3.6.4070	First release
Role	admin	
Example	`switch (config)# ipv6 dhcp relay instance 1 interface-id option`	
Related Commands		
Note		

ipv6 dhcp relay instance vrf

ipv6 dhcp relay instance <instance-id> vrf <vrf-name>
no ipv6 dhcp relay instance <instance-id> vrf <vrf-name>

Configures instance in the given VRF.
The no form of the command will reset the instance back to default VRF.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	vrf-name	Name of VRF
Default	Default VRF	
Configuration Mode	config	
History	3.6.4070	First release
Role	admin	
Example	`switch (config)# ipv6 dhcp relay 1 vrf test`	
Related Commands		
Note	When an instance is moved from one VRF to another - it loses all its current configuration.	

ipv6 dhcp relay instance port

ipv6 dhcp relay instance <instance-id> port <udp-port>
no ipv6 dhcp relay instance <instance-id> port <udp-port>

Modifies the UDP port for the given instance.
The no form of the command will set the UDP port to default value.

Syntax Description	instance-id	DHCP instance ID. Range: 1-8.
	port	UDP Port ID Range: 1-65534 Default: 547
Default	UDP port 547	
Configuration Mode	config	
History	3.6.4070	First release
Role	admin	
Example	`switch (config)# ipv6 dhcp relay 1 port 555`	
Related Commands		
Note		

ipv6 dhcp relay instance interface-id option

ipv6 dhcp relay instance <instance-id> interface-id option [user-defined-id]

Specifies the content of the interface-id option that will be sent by the relay agent.

Syntax Description	instance-id	DHCP instance ID. Range: 1-8.
	user-defined-id	Interface-id option content. Length: 1-15 (char) Default: interface name
Default	N/A	
Configuration Mode	config	
History	3.6.4070	First release
Role	admin	
Example	`switch (config)# ipv6 dhcp relay instance <instance-id> interface-id option eth1/1`	
Related Commands		
Note		

ipv6 dhcp relay instance use-secondary-ip

	ipv6 dhcp relay instance use-secondary-ip **no ipv6 dhcp relay instance use-secondary-ip** Enables the switch to relay a single request from the client multiple times simultaneously, with each of the IP addresses configured on the corresponding downstream interfaces as the respective gateway address (giaddr field of IPv6 DHCP request packet). The no form of the command disables this function.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.8008
Role	admin
Example	`switch (config ipv6 dhcp relay instance 1)# use-secondary-ip`
Related Commands	N/A
Note	

clear ipv6 dhcp relay counters

clear ipv6 dhcp relay counters [vrf {<vrf-name> | all} | instance <instance-id>]

Clears DHCP relay counters for specific instance or all instances in given VRF or all instances in the system.

Syntax Description	vrf-name	VRF name or “all” for all VRFs
	instance-id	DHCP instance ID Range: 1-8
Default	N/A	
Configuration Mode	config	
History	3.6.4070	First release
	3.6.5000	Added “all” parameter
Role	admin	
Example	`switch (config)# clear ipv6 dhcp relay counters vrf all`	
Related Commands		
Note		

6.10.4.1.2Show

show ipv6 dhcp relay

show ipv6 dhcp relay [instance <instance-id>]

Displays general DHCP configuration on all instances.
If instance id is defined then specific instance configuration is displayed.

Syntax Description	instance-id	DHCP instance ID. Range: 1-8.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4070	First release
	3.6.5000	Updated Example
	3.6.6000	Updated Example
	3.6.8008	Updated Example
Role	admin	
Example		

```
switch (config)# show ipv6 dhcp relay

Instance ID 1:
  VRF Name: default

  DHCP Servers:
    2001:db8:701f::8f9

  DHCP relay agent options:
    All_DHCP_Servers   : Disabled
    Interface-id Option: Disabled
    UDP port           : 547
    Auto-helper        : Disabled
    Status             : Down

  ----------------------------------------
  Interface   Label             Mode
  ----------------------------------------
  eth1/5      N/A               downstream
```

Related Commands	
Note	• If no DHCP-R instance is given, then all DHCP-R instances are displayed • Only configured interfaces are displayed • Once vrf-auto-helper is enabled, no interface is displayed

show ipv6 dhcp relay counters

show ipv6 dhcp relay counters [instance <instance-id> | vrf <vrf-name>]

Displays the DHCPv6 relay counters.

Syntax Description	instance-id	Displays the DHCPv6 relay counters for a given instance
	vrf	Displays the DHCPv6 relay counters in a given VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.8008	Updated Example
Role	admin	

Example

```
switch (config) # show ipv6 dhcp relay counters

Instance 1:
  VRF Name: vrf-default

  DHCP Counter flags:
    SPR : Server Packets Received
    SPE : Server Packets Error
    SPRE: Server Packet Relayed
    CPR : Client Packets Received
    RP  : Relay Packets
    RE  : Relay Errors

  ---------------------------------
  Req/Resp    Received    Forwarded
  ---------------------------------
  All Req     0           0
  All Res     0           0

  -----------------------------------------------------
  If         SPRE       SPE        SPR        CPR
  -----------------------------------------------------
  eth1/5     0          0          0          0

  Packets Relayed to Server:
    ---------------------------------------------------------------------
    Server                                          RP          RE
    ---------------------------------------------------------------------
    2001:db8:701f::8f9                              0           0
```

Related Commands	N/A
Note	

Appendix A: Enhancing System Security According to NIST SP 800-131A

Mellanox switch systems comply by default with NIST SP 800-131A as described in the table below.

Table 61 - Supported Event Notifications and MIB Mapping

Component	Configuration	Command
HTTP	HTTP disabled	no web http enable
HTTPS	HTTPS enabled	no web https enable
	SSL ciphers = TLS1.2	web https ssl ciphers all
	SSL renegotiation disabled	web https ssl renegotiation enable
SSH	SSH version = 2	ssh server min-version 1
	SSH ciphers = aes256-ctr, aes192-ctr, aes128-ctr, aes128-gcm@openssh.com, aes256-gcm@openssh.com	no ssh server security strict

A.1 Overview

This appendix describes how to enhance the security of a system in order to comply with the NIST SP 800-131A standard. This standard is a document which defines cryptographically "acceptable" technologies. This document explains how to protect against possible cryptographic vulnerabilities in the system by using secure methods. Because of compatibility issues, this security state is not the default of the system and it should be manually set.

Some protocols, however, cannot be operated in a manner that complies with the NIST SP 800-131A standard.

A.2 Web Certificate

Onyx supports signature generation of sha256WithRSAEncryption, sha1WithRSAEncryption self-signed certificates, and importing certificates as text in PEM format.

➢ ***To configure a default certificate:***

Step 1. Create a new sha256 certificate. Run:

```
switch (config) # crypto certificate name <cert name> generate self-signed hash-algo-
rithm sha256
```

For more details and parameters refer to the command crypto certificate name in the *Mellanox Onyx User Manual*.

Step 2. Show crypto certificate detail. Run:

```
switch (config) # show crypto certificate detail
```

Search for “signature algorithm” in the output.

Step 3. Set this certificate as the default certificate. Run:

```
switch (config) # crypto certificate default-cert name <cert name>
```

➢ ***To configure default parameters and create a new certificate:***

Step 1. Define the default hash algorithm. Run:

```
switch (config) # crypto certificate generation default hash-algorithm sha256
```

Step 2. Generate a new certificate with default values. Run:

```
switch (config) # crypto certificate name <cert name> generate self-signed
```

When no options are selected, the generated certificate uses the default values for each field.

To test strict mode connect to the WebUI using HTTPS and get the certificate. Search for “signature algorithm”.

There are other ways to configure the certificate to sha256. For example, it is possible to use `certificate generation default hash-algorithm` and then regenerate the certificate using these default values. Please refer to the *Mellanox Onyx User Manual* for further details.

It is recommended to delete browsing data and previous certificates before retrying to connect to the WebUI.

Make sure not to confuse “signature algorithm” with “Thumbprint algorithm”.

A.3 Code Signing

Code signing is used to verify that the data in the image is not modified by any third-party. MLNX-OS supports signing the image files with SHA256, RSA2048 using GnuPG.

Strict mode is operational by default.

A.4 SNMP

SNMPv3 supports configuring username, authentication keys and privacy keys. For authentication keys it is possible to use MD5 or SHA. For privacy keys AES or DES are to be used.

- ***To configure strict mode, create a new user with HMAC-SHA1-96 and AES-128. Run:***

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128
<password2>
```

- ***To verify the user in the CLI, run:***

```
switch (config) # show snmp user
```

To test strict mode, configure users and check them using the CLI, then run an SNMP request with the new users.

For more information please refer to the *Mellanox Onyx User Manual*.

SNMPv1 and SNMPv2 are not considered to be secure. To run in strict mode, only use SNMPv3.

A.5 SSH

The SSH server on the switch by default uses secure ciphers only, message authentication code (MAC), key exchange methods, and public key algorithm. When configuring SSH server to strict mode, the aforementioned security methods only use approved algorithms as detailed in the NIST 800-181A specification and the user can connect to the switch via SSH in strict mode only.

➢ ***To enable strict security mode, run:***

```
switch (config) # ssh server security strict
```

The following ciphers are disabled for SSH when strict security is enabled:

- 3des-cbc
- aes256-cbc
- aes192-cbc
- aes128-cbc
- arcfour
- blowfish-cbc
- cast128-cbc
- rijndael-cbc@lysator.liu.se

The no form of the command disables strict security mode.

Make sure to configure the SSH server to work with minimum version 2 since 1 is vulnerable to security breaches.

➢ ***To configure min-version to strict mode, run:***

```
switch (config) # ssh server min-version 2
```

Once this is done, the user cannot revert back to minimum version 1.

A.6 HTTPS

By default, Onyx supports HTTPS encryption using TLS1.2 only. Working in TLS1.2 mode also bans MD5 ciphers which are not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- RSA_WITH_AES_128_CBC_SHA256
- RSA_WITH_AES_256_CBC_SHA256
- DHE_RSA_WITH_AES_128_CBC_SHA256
- DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384

- ***To enable all encryption methods, run:***

```
switch (config) # web https ssl ciphers all
```

- ***To enable only TLS ciphers (enabled by default), run:***

```
switch (config) # web https ssl ciphers TLS
```

- ***To enable HTTPS strict mode, run:***

```
switch (config) # web https ssl ciphers TLS1.2
```

- ***To verify which encryption methods are used, run:***

```
switch (config)# show web
Web User Interface:
 Web interface enabled: yes
 HTTP enabled: yes
 HTTP port: 80
 HTTP redirect to HTTPS: no
 HTTPS enabled: yes
 HTTPS port: 443
 HTTPS ssl-ciphers: TLS1.2
 HTTPS certificate name: default-cert
 Listen enabled: yes
 No Listen Interfaces.

 Inactivity timeout: disabled
 Session timeout: 2 hr 30 min
 Session renewal: 30 min

Web file transfer proxy:
 Proxy enabled: no

Web file transfer certificate authority:
 HTTPS server cert verify: yes
 HTTPS supplemental CA list: default-ca-list
switch (config)#
```

On top of enabling HTTPS, to prevent security breaches HTTP must be disabled.

- ***To disable HTTP, run:***

```
switch (config) # no web http enable
```

A.7 LDAP

By default, supports LDAP encryption SSL version 3 or TLS1.0 up to TLS1.2. The only banned algorithm is MD5 which is not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- DHE-DSS-AES128-SHA256
- DHE-RSA-AES128-SHA256
- DHE-DSS-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256

- DHE-DSS-AES256-SHA256
- DHE-RSA-AES256-SHA256
- DHE-DSS-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- ECDH-ECDSA-AES128-SHA256
- ECDH-RSA-AES128-SHA256
- ECDH-ECDSA-AES128-GCM-SHA256
- ECDH-RSA-AES128-GCM-SHA256
- ECDH-ECDSA-AES256-SHA384
- ECDH-RSA-AES256-SHA384
- ECDH-ECDSA-AES256-GCM-SHA384
- ECDH-RSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- AES128-SHA256
- AES128-GCM-SHA256
- AES256-SHA256
- AES256-GCM-SHA384

➢ ***To enable LDAP strict mode, run:***

```
switch (config) # ldap ssl mode {start-tls | ssl}
```

Both modes operate using SSL. The different lies in the connection initialization and the port used.

Appendix B: Feature Support per IC and CPU Type

Table 62 lists which features are supported by which IC family and CPU type.

New features added on release 3.6.81xx and beyond are supported on Spectrum-based switches only.

Table 62 - Feature Support (Y for Supported, N for Not Supported)

Feature	SwitchX®-2 PPC	SwitchX®-2 x86	Spectrum™ Family
Image Maintenance via Mellanox ONIE	N	Y	Y
IPv6	N	N	Y
JSON	N	Y	Y
OpenFlow 1.0	Y	Y	Y
OpenFlow 1.3	N	N	Y
PIM	N	N	Y
PTP	N	N	Y
QoS RED & ECN	N	N	Y
S&F config	N	N	Y
Signal Degradation Monitoring	N	N	Y
Shared Buffers	N	N	Y
Storm Control	N	N	Y
Telemetry (histograms and threshold)	N	N	Y
User Defined Keys	N	N	Y
VXLAN	N	N	Y

Appendix C: Splunk Integration with Mellanox Products

Splunk automatically clusters millions of log records in real time back into their patterns and finds connections between those patterns to form the baseline flows of each software individually, thus enables you to search, monitor and analyze that data to discover powerful insights across multiple use cases.

This appendix provides a guide on the first steps with Splunk and helps you to begin enjoying reduced time in detecting and resolving production problems.

C.1 Getting Started with Splunk

Step 1. Download Splunk and extract the Splunk Enterprise version. (Splunk software is available as an RPM or TGZ.)

Step 2. Create a Splunk User /group. Run:

```
[root@server] groupadd splunk
[root@server] useradd -d /opt/splunk -m -g splunk splunk
```

Step 3. Splunk installation. Run:

```
[root@server] tar -xzvf splunk-7.0.0-c8a78efdd40f-Linux-x86_64.tgz
[root@server] ls
```

Step 4. A new folder called Splunk is created.

```
[root@server] cp -rp splunk/* /opt/splunk/
[root@server] chown -R splunk: /opt/splunk/
[root@server] su - splunk
[splunk@server] cd bin
[splunk@server] ./splunk start --accept-license
```

Now you can access your Splunk WebUI at http://IP:8000/ or http://hostname:8000/. You need to make sure that port 8000 is open in your server firewall.

C.2 Switch Configuration

In this example we are not using the default UDP port 514 to show that any other port can be also used.

Step 5. In order to add a task, the switch must be configured to send logs to our Splunk server. Run:

```
switch > enable
switch # configure terminal
switch (config) # show snmp
SNMP enabled:         yes
SNMP port:            161
System contact:
System location:

Read-only communities:
   public
```

```
Read-write communities:
   (none)

Interface listen enabled: yes
No Listen Interfaces.

switch (config) # snmp-server host 10.212.23.1 informs port 8597
switch (config) # snmp-server host 10.212.23.1 traps port 8597
switch (config) # snmp host 10.212.23.1 informs 8597
switch (config) # snmp host 10.212.23.1 traps  8597


Summary configuration:


switch (config) # show running-config
## Logging configuration
##
   logging 10.212.23.1
   logging 10.212.23.1 port 8597
   logging 10.212.23.1 trap info
   logging 10.212.23.1 trap override class events priority err
   logging monitor events notice
   logging receive
## SNMP configuration
no snmp-server host 10.209.21.221 disable
snmp-server host 10.209.21.221 traps port 8597 version 2c
no snmp-server host 10.212.23.1 disable
snmp-server host 10.212.23.1 traps port 8597 version 2c 8597
```

C.3 Adding a Task

Step 6. The first screen encountered after signing into the Splunk WebUI includes the “Add Data” icon.

Figure 42: Add Data Option

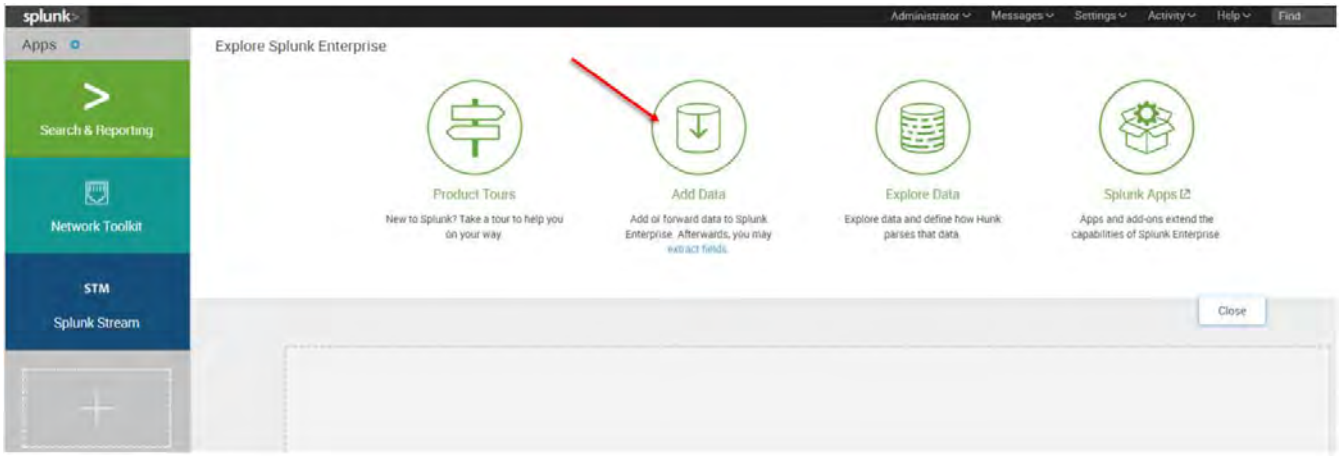

Step 7. The “Add Data” tab opens up with three options: Upload, Monitor, and Forward. Here our task is to monitor a folder, so we click Monitor. to proceed

Figure 43: Monitor Icon

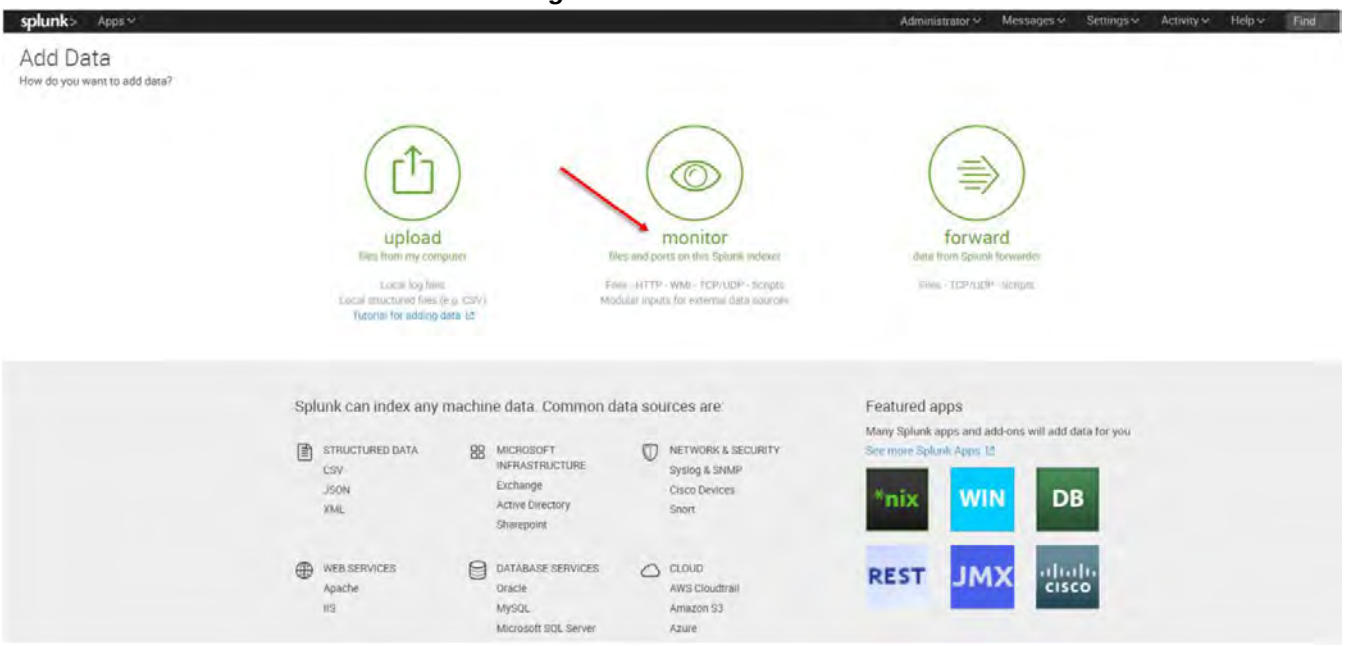

In the Monitor option, the following four categories are available:

- File & Directories: Monitor files/folders
- HTTP Event Collector: Monitor data streams over HTTP
- TCP/UDP: Monitor service ports
- Scripts: Monitor scripts

C.4 Retrieving Data from TCP and UDP Ports

Step 8. Per our current purpose, we choose TCP/UDP option.

Figure 44: TCP/UDP

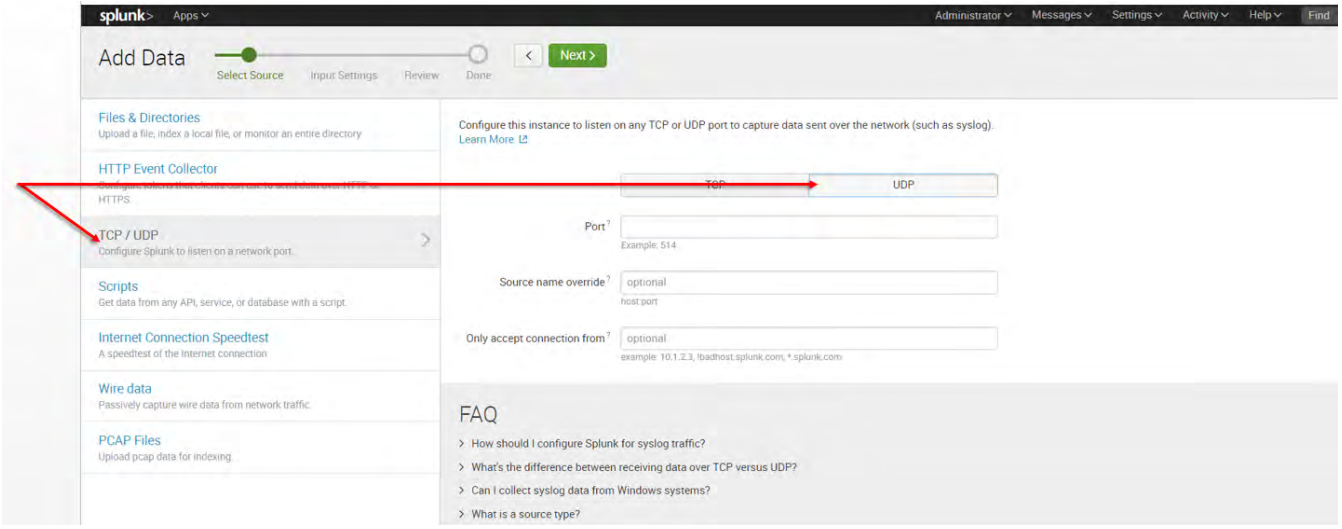

Step 9. Click the TCP or UDP button to choose between a TCP or UDP input, and enter a port number in the "Port" field.

Step 10. In the "Source name override" field, enter a new source name to override the default source value, if required.

Figure 45: TCP/UDP Fields

Step 11. Click “Next” to continue to the Input Settings page where we will create a new source type called Mellanox-Switch.

Figure 46: Input Settings

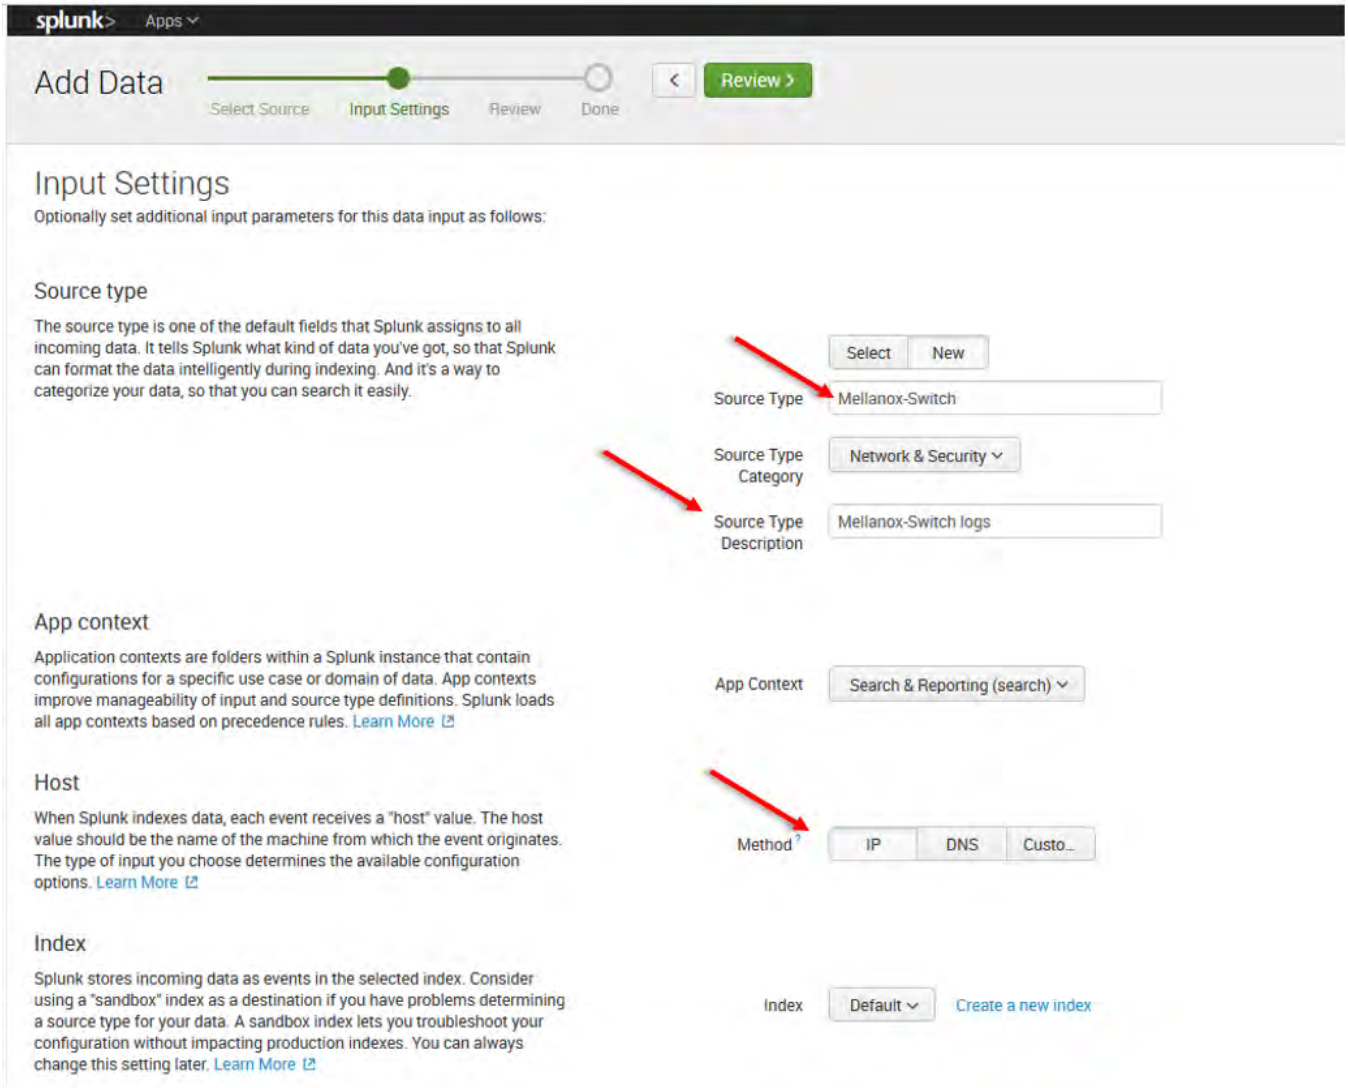

Step 12. Click Next > Review > Done > Start Searching

Figure 47: Start Searching

C.5 SNMP Input to Poll Attribute Values and Catch Traps

SNMP represents an incredibly rich source of data that you can get into Splunk for visibility across a very diverse IT landscape.

SNMP agents may also send notifications, called Traps, to an SNMP trap listening daemon.

C.5.1 Getting Started

Browse to Splunkbase and download the SNMP Modular Input from https://splunkbase.splunk.com/app/1537/.

To install, simply untar the file to SPLUNK_HOME/etc/apps and restart Splunk.

C.5.2 Configuration

Login to the Splunk WebUI and go to Manager > Add Data > Monitor > SNMP > New, and set up your input data.

Figure 48: SNMP

Figure 49: SNMP Attributes Polling Settings

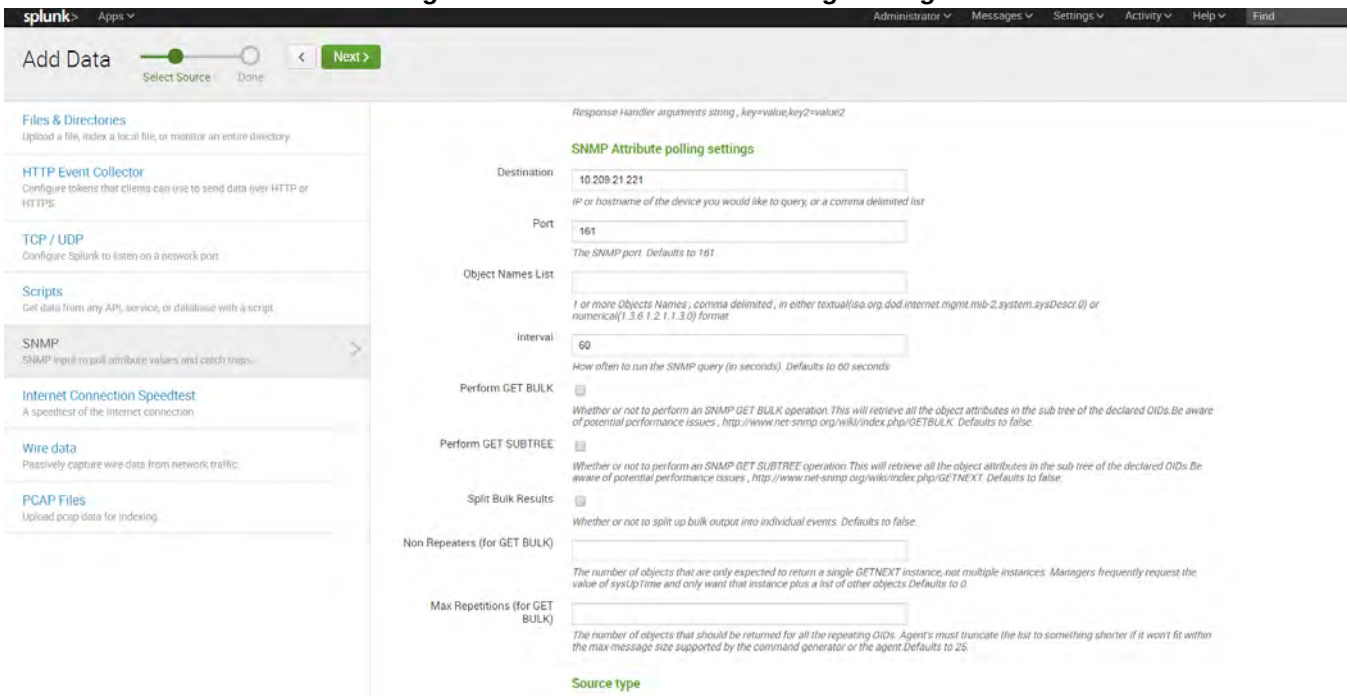

Figure 50: SNMP Attributes Polling Settings

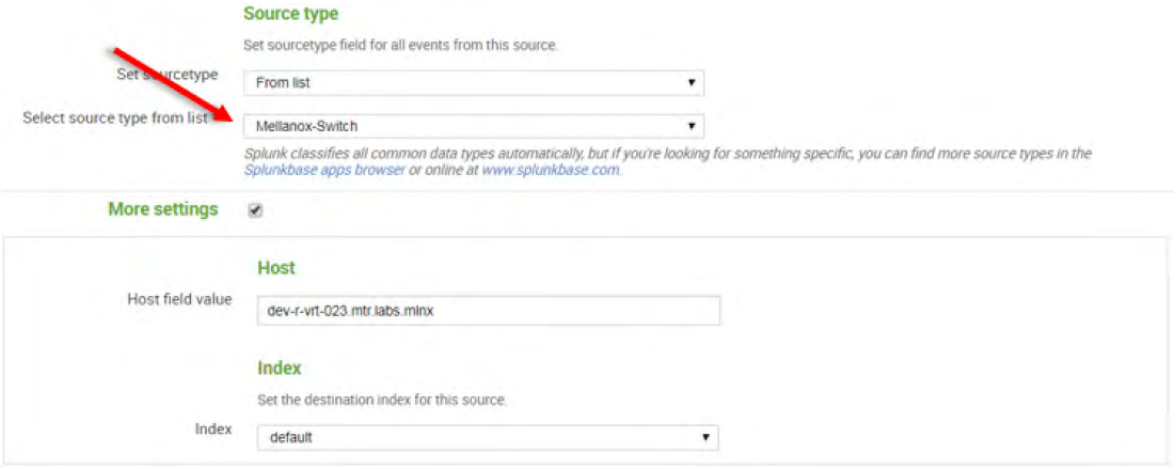

Step 13. After configuration is complete it is recommend to run Mellanox-Switch again: Search > Data Summary > Sourcetypes > Mellanox-Switch.

Figure 51: Mellanox-Switch

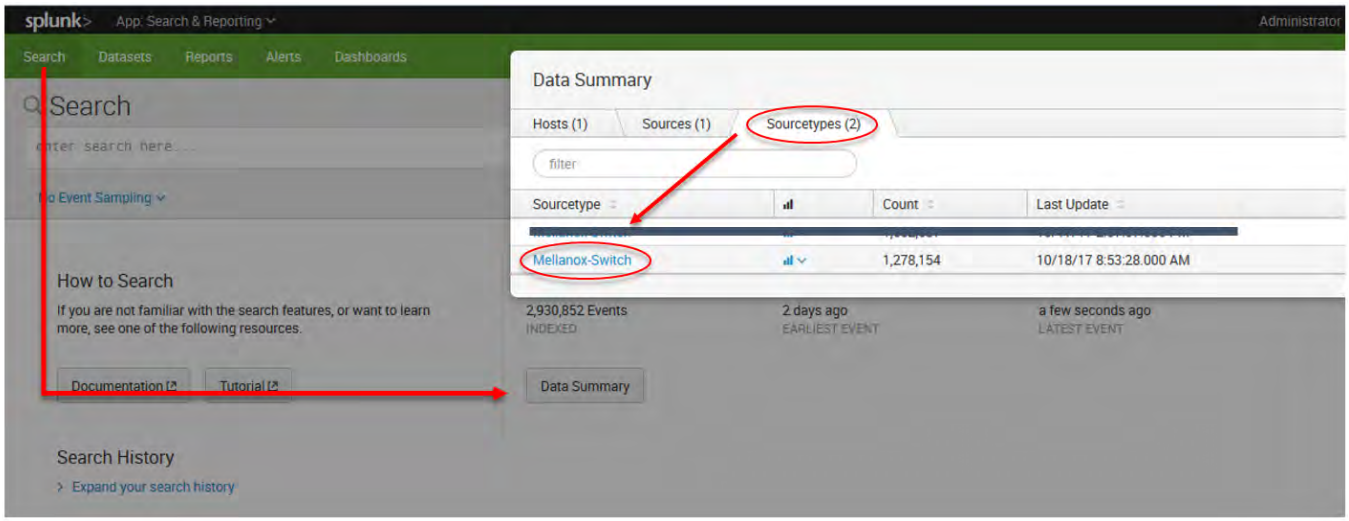

Step 14. Select “Mellanox-Switch” and “Add to search”.

Figure 52: Add to Search

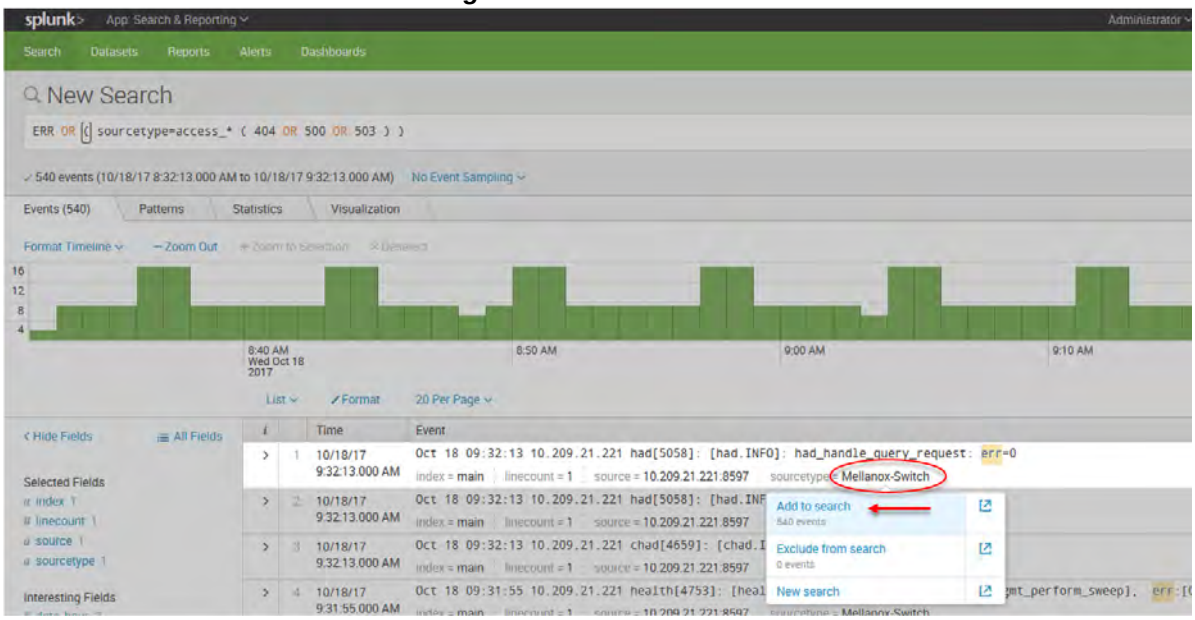

Step 15. You can add to search any value that is relevant for you.

Figure 53: Search Options

Patterns can be viewed not on real time and you can create alert on most repeatable events.