



PicOS Routing and Switching Configuration Guide

October, 2017

Version: 1

www.pica8.com

Pica8, Inc.
1032 Elwell Court, Suite 105
Palo Alto, CA. 94303
+1 (650) 614-5838

sales@pica8.com
support@pica8.com

Contents

Supported Features	25
Command-Line Interface	26
System Administration	27
Ethernet Ports Configuration	28
Layer 2 Switching Configuration	29
Layer 3 Unicast Routing Configuration	31
Layer 3 Multicast Routing Configuration	33
QoS Configuration	34
OpenFlow in Crossflow Mode	35
VXLAN Configuration	36
L2/L3 Troubleshooting Guide	37
Configuration Appendix	38
Summary of Supported Services	39
Command-Line Interface	41
CLI Configuration	41

From Linux Shell to L2/L3 Shell	42
Operation Mode and Configuration Mode	42
Displaying the Current Configuration	43
Display Setting Configuration	43
Rolling Back a Configuration	44
Managing Configuration Files	44
Saving and Loading Configuration Files	48
Commit Confirmed	49
Commit Check	49
Commit Failed and Exit Discard	51
Configuring a Command Alias	51
Configure L2/L3 from Linux Shell	51
Bash Linux Shell	52
PicOS Upgrade and Configuration Change	53
Set CLI	54
System Administration	55
Configuring DHCP and Static IP Address	56
Enabling DHCP	56
Configuring a Static IPv4 Address	56
Configuring a Static IPv6 Address	56
Enabling DHCP Relay in a VLAN Interface	56
Configuring DHCP Relay after PicOS 2.9.2(including 2.9.2)	57
Procedure	57
Configuration Example	58
Configuring DHCP Option82	60
Enable DHCP Option82	60
Modify the Circuit-id of Option82	60
Configuring DHCP Snooping	60
Enable DHCP Snooping	61
Configure DHCP Snooping Binding File and Timeout	61
Configure DHCP Snooping Trust Port	61
Display the DHCP Snooping Table of Host Information	61
Configuring DHCP Snooping after PicOS 2.9.2(including 2.9.2)	61

Procedure	62
Configuration example	62
Configuring DHCP Relay and DHCP Snooping together after PicOS 2.9.2(including 2.9.2)	63
Configuration example	64
Configuring a User Account	65
Creating a User Class and Password	66
Configuring a T elnet Announcement	66
Configuring Authentication, Authorization, and Accounting	66
Configuring AAA in the Switch	66
Displaying AAA information	67
Configuring the Local Log-in	69
Configuring SSH and Telnet Parameters	69
Configuring the SSH Connection Limit	69
Disabling Telnet Service	70
Enabling and Disabling Inband Service	70
Configuring the Log-in ACL	70
Configuring NTP and the Time Zone Parameter	70
Configuring the NTP Server IP Address	70
Configuring the Time Zone	71
Configuring the S ystem Clock	71
Configuring PTP	71
Principle	71
Introduction to PTP	71
Basic Concepts	72
Delay Request-Response Mechanism	74
Product Support	75
PTP Configuration	76
Procedure	76
Configuration Example	76
Configuring the linux-config-unreliable	77
Xorp Control	77
Bash Control	78
Configuring IPFIX Parameters	78
Configuring sFlow	78
Globally Enabling sFlow	79
Configuring sFlow Parameters	79
Configuring sFlow on a Specific Interface	79
Configuring SNMPv2	80
Configuring SNMPv2 Parameters	80

Configuring an SNMP ACL	81
Configuring SNMPset	81
Configuring the Syslog Log Level	82
Configuring the Syslog Level	82
Configuring the SNMP Logging Facility	82
Configuring the Syslog Disk and Syslog Host	83
Configuring the Syslog Host	83
Configuring Syslog for Local Storage	83
Displaying System Information	83
Displaying the System Fan	84
Displaying the System Power Supply Unit	84
Displaying the System Serial Number	84
Displaying Additional System Information	84
Technical Support	87
Flushing ARP and the Neighbor Table	88
Flushing the ARP Entry	88
Rebooting the System	88
Displaying the Debugging Message	88
Syslog Monitor On	89
System Management Command List	89
IPv6 Management Support	89
SNMP over IPv6	89
Telnet/SSH over IPv6	89
Syslog over IPv6	89
NTP Client IPv6	90
DNS Client IPv6	90
DHCPv6 Client	90
Radius over IPv6	90
PoE Configuration	90
PoE Configuration	90
Show Help Information	91
Show Version Information	91
Key Words	91
PSE Function Enable	91
Powerup Command	92
Logical Port Map Enable	92
Reset the Specified Port or Ports	93
Reset the Error Statistics on Specified Ports	93
Set the PSE Functionality on all Ports	94

Set the Default Maximum Allocated Power Parameters in High Power Mode on all Ports in the PoE Subsystem	94
The Switch Port from Low Power Mode to High Power Mode	95
Reset the Entire PoE Subsystem	95
Global System Parameters Configuration Command	95
Configure Detection Type on Specified Ports	96
Configure the Classify Type on Specific Ports	96
Configure Auto Mode on Specific Ports	97
Configure Disconnect Type on Specific Ports	97
Configure Power Threshold Type on Specified Ports	98
Configure Max Threshold on Specific Ports	98
Configure Power Management Mode	98
Configure Power Supply Details for Each Combination of the MPSS Pins	99
Configure Port Priority on Specific Ports	99
Configures Port Power Up Mode	100
Configure the Per Port Power Disconnect Mask	100
Acquire Port Status	100
Acquires the Total Power Allocated	101
NETCONF Configuration	101
Configuring SNMPv3	103
	104
Configuring Basic Information	104
Configuring trap-group	105
Setting Up a User	105
Configuring Mib-view	105
Configuring Security-level	106
NMS Visits Switch by user	106
Ethernet Ports Configuration	107
Physical Ethernet Port Configuration	108
Configuring the MTU and Rate-limit	108
Enabling Port Flow Control	108
Split 40GE Ports into 4x10GE Ports	108
Configuring Port Speed	110
VLAN Port Configuration	110
Configure the Native VLAN ID	110
Add a Port to a VLAN with Tagged Packet	111
Add a Port to a VLAN with Untagged Packet	111
Creating a VLAN within the VLAN Range	112
VLAN Configuration Example	112
Configuring Switch A	113
Configuring Switch B	114

40G Changes to 4*10G in L2/L3	115
40G Changes to 4*10G in L2/L3 mode on P-5101	117
In L2/L3 mode configure	117
QSFP (8 x 40G+40*10G)	117
SFP (72x 10G)	119
40G Changes to 4*10G in L2/L3 mode on P-5401	122
In L2/L3 mode configure	122
QSFP (32 x 40G)	122
SFP-64 (16 x 40G + 64 x 10G)	123
SFP (8 x 40G + 96 x 10G)	127
40G Changes to 4*10G in L2/L3 mode on accton_as6701_32x	131
In L2/L3 mode configure	131
SFP (8 x 40G + 96 x 10G)	133
40G Changes to 4*10G in L2/L3 mode on P-3922	137
In L2/L3 mode configure	137
QSFP(48*10G+4*40G)	137
SFP (64*10G)	139
40G Changes to 4*10G in L2/L3 mode on P-3920	142
In L2/L3 mode configuration	142
QSFP(48*10G+4*40G)	142
SFP (64*10G)	144
40G Changes to 4*10G in L2/L3 mode on accton_as5712_54x/HP5712	146
In L2/L3 mode configuration	146
QSFP(48*10G+6*40G)	147
40G Changes to 4*10G in L2/L3 mode on P-3930	151
In L2/L3 mode configure	151
QSFP(48*10G+4*40G)	152
40G Changes to 4*10G in L2/L3 mode on Niagara2632XL	156
In L2/L3 mode configuration	156
QSFP (32 x 40G)	156
SFP-64 (16 x 40G + 64 x 10G)	158
SFP (8 x 40G + 96 x 10G)	161
In L2/L3 mode configuration	165
40G Changes to 4*10G in L2/L3 mode on as5812_54t	170
In L2/L3 mode configure	170
40G Changes to 4*10G in L2/L3 mode on arctica4806xp	175
In L2/L3 mode configure	175
QSFP(48*10G+6*40G)	176
LLDP Configuration (Link Layer Discovery Protocol)	180
Configuring the LLDP Mode	181
Selecting Optional TLVs	181
Displaying LLDP Information	181
Configuring Other Parameters	182
Voice Vlan with LLDP Configuration Guide	182

Basic Configuration:	182
Voice Vlan with LLDP Compliance CDP Configuration Guide	184
Basic Configuration:	185
Voice Vlan Configuration Guide	188
Default OUI:	188
Auto Mode:	189
Manual Mode:	189
Basic Configuration	189
LLDP MED Configuration	190
Basic Configuration:	190
Voice Vlan with LLDP MED Configuration Guide	192
Basic Configuration:	193
Static Link Aggregation (LAG) Configuration	195
Displaying S tatic LAG Information	195
LAG Hashing Configuration	196
LAG Hashing Configuration and Example	196
LAG Hash Mapping	196
Resilient LAG Hashing Configuration and Example	197
Link Aggregation Control Protocol (LACP) Configuration	197
Configuring LACP LAGs	198
Displaying LACP LAG Information	198
Storm Control in Ethernet Port Configuration	198
Configuring Storm Control in packet per second on ingress port	199
Configuring Storm Control in ratio on ingress port	199
Configuring Storm Control in kilobits per second on ingress port	199
UDLD Configuration	199
Configure UDLD normal mode on global	200
Configure UDLD aggressive mode on global	200
Configure UDLD normal mode on Specific Port	200
Configure UDLD aggressive mode on Specific Port	200
Configure UDLD Message-interval	200
UDLD status as follows when the link is good	200
UDLD status as follows when UDLD detect uni-directional fault	201
UDLD status when self loop link	201
Recover port status Disabled(UDLD)	202
Port Security Configuration	203
Enabling Port Security	203
Configuring the Maximum Number of Secure Dynamically Learned MAC Addresses	203
Configuring Static Secure MAC Addresses on a Port	204
Configuring Port Security with Sticky MAC Addresses on a Port	204

Configuring Secure MAC Address Aging Time	204
Configuring Port Security Violation Mode on a Port	205
Configuring Port Security Auto-recovery Time	205
Recovering the Port in Error-discard	205
Configuring Port Security Block Mode on a Port	206
Displaying Port Security Settings	206
Disabling Port Security	207
802.1x Configuration	207
Enabling/ disabling 802.1x on the Interface	208
Configuring 802.1x Port Control Mode of Interface	208
Configuring the Value of Reauth-period	208
Deleting Protocols of dot1x	209
Buffer management	210
Introduction	210
Command configuration	211
1.multicast queue	211
2.unicast queue	211
3.dynamic threshold ratio configuration	212
Application notice	213
Layer 2 Switching Configuration	214
Static MAC entries and Dynamic MAC Address Learning	215
Configuring a Static MAC Entry and Managing the FDB	215
Cut-Through Switching Method	215
Configuring your Switch to the Store-and-Forward Method	215
MLAG Configuration Guide	216
Configuring MLAG domain-id	217
Configuring MLAG Peer	217
Configuring MLAG Priority	217
Configuring MLAG hello-interval	217
Configuring a Basic MLAG Step-by-Step Procedure	219
Configuring a Basic MLAG Example	220
Configuring Switch A with Static and LACP LAG	220
Configuring an Aggregation Interface to VLAN Members	221
Configure the L3 Interface IP Address	221
Configuring the domain-id for the MLAG	221
Configuring the Peer IP Address and the Peer-link for the MLAG Domain Peer	222
Configuring Switch B with Static and LACP LAG	222
Configuring an Aggregation Interface to VLAN Members	222
Configure the L3 Interface IP Address	223
Configuring the domain-id for the MLAG Domain	223
Configuring the Peer IP Address and the peer-link for the MLAG Domain Peer	223

Configuring Switch C with LACP and LAG	224
Configuring an Aggregation Interface to VLAN Members	224
View the MLAG Internal and Neighbor Status	224
View the MLAG Internal and Neighbor Status of Switch B	225
If the configuration is inconsistent, the mlag internal as follows:	225
Configuring a MLAG domain with MSTP example	225
Configuring Switch A with LACP LAG	226
Configuring an Aggregation Interface to VLAN Members	227
Configure the L3 interface IP address	227
Configuring the domain-id for the MLAG	227
Configuring the peer IP address and the peer-link for the MLAG domain peer	228
Configuring Switch B with LACP LAG	228
Configuring an Aggregation Interface to VLAN Members	228
Configure the L3 interface IP address	229
Configuring the domain-id for the MLAG domain	229
Configuring the peer IP address and the peer-link for the MLAG domain peer	229
Configuring Switch C and Switch D with LACP LAG	229
Configuring Switch C and Switch D an aggregation interface add to VLAN Members	230
View the MLAG internal and neighbor status of Switch A	230
View the MSTP status of Switch A	230
View the MLAG internal and neighbor status of Switch B	231
View the MSTP status of Switch B	231
Configure Mlag and Active-Active-Vrrp example	231
Requirements	231
Overview	231
Topology	232
Configure on Switch A	233
Configure Mlag on Switch A	233
Configure Vrrp on Switch A	233
Configure on Switch B	233
Configure Mlag on Switch B	233
Configure Vrrp on Switch B	234
Configure on Switch C	234
Configure Mlag on Switch C	234
Configure Vrrp on Switch C	234
Configure on Switch D	234
Configure Mlag on Switch D	234
Configure Vrrp on Switch D	235
Configuring a MLAG with vxlan Example	235
01 Configuring mlag in Switch A and Switch B	236
Switch A Configuration	236
Switch B Configuration	237
Configuring Lag in Switch C and Switch D	239
Switch C and Switch D Configuration	239
Configure the vxlan in Switch A and Switch B	240

Q-in-Q Basic Port Configuration	240
Configuring the Q-in-Q Tunneling Internal/External Mode	240
Configuring Q-in-Q Tunneling to Map Ingress VLANs to Service VLANs	241
Configuring Q-in-Q Tunneling Egress Pop Service VLANs	242
Q-in-Q Configuration Example	243
Configuration on Provider A	243
Configuration on Provider B	246
MSTP Configuration	248
Enabling Spanning Tree Mode in MSTP	248
Configuring Basic Global Parameters of MSTP	248
Configuring MSTP Interface Parameters	250
Configuring the BPDU Filter	250
Configuring BPDU Root Guard	250
Configuring BPDU TCN-Guard	250
Configuring MSTP BPDU-Guard	251
Disabling/Enabling MSTP	252
MSTP Configuration Example	252
Configuring Switch A, Example 1	254
Configuring Switch B, Example 1	255
Configuring Switch C, Example 1	256
Configuring Switch D, Example 1	256
Configuring Switch E, Example 1	257
Configuring Switch A, Example 2	260
Configuring Switch B, Example 2	261
Configuring Switch C, Example 2	262
Configuring Switch D, Example 2	264
Configuring Switch E, Example 2	265
PVST Configuration	266
Enabling Spanning Tree Mode in PVST	266
Configuring Basic VLAN Parameters of PVST	266
Configuring PVST Interface Parameters	266
Configuring the Interface Mode	267
Configuring PVST BPDU-guard	267
Disabling/Enabling PVST on One VLAN	268
Disabling/Enabling PVST	269
PVST Configuration Example	270
Configuring Switch A	270
Configuring Switch B	271
Configuring Switch C	272
Configuring Switch D	273
Configuring Mirroring	273
Configuring Mirroring to Analyze Traffic	274

Configuring Mirroring Guide	274
Configuring a Port as a Mirroring Port	274
Configuring Mirroring on Egress Port or Ingress Port	274
Configure Monitor the flows of ingress port:	274
Configure Monitor the Flows of Egress Port:	274
Mirroring Configuration Example	275
Configure Mirroring Port	275
Configuring Mirroring on Egress or Ingress Port	276
Generate Traffic	276
Note:	276
BPDU Tunneling Configuration	276
Configuring BPDU Tunneling for STP on an Interface	276
Configuring Destination Multicast MAC Address for BPDU Packets	277
BPDU Tunneling Configuration Example	277
Configuration on Provider A	277
Configuration on Provider B	278
Configuring Flex Links Preemption Delay	279
Configuring the Preemption Mode	280
Showing Flex Links on All Interfaces	280
Configuring IPv6 RA Guard	280
Configuring Trusted-Port	281
Displaying RA Guards	281
LFS Configuration	282
Abstract	282
Link Fault Signaling	282
LFS Commands	282
Up Mode	283
Up Mode Commands	283
Symmetric Hash for ECMP and LAG Configuration Example	283
Layer 3 Unicast Routing Configuration	286
Layer 3 VLAN Interface Configuration	288
ARP Configuration	289
Configuring ARP Aging Time	289
Configuring a Static ARP Entry	289
Dynamic ARP Inspection	289
Default Administrative Distance Values	291
Select the Best Path	291
Default Distance Value Table	291
Static Routing Configuration	292

Configuration Example	294
Configuring Switch A	294
Configuring Switch B	295
Configuring Switch C	296
RIPv2 Routing Protocol Configuration	297
RIPv2 Routing Configuration Example	298
Configuration Example	298
Configuring Switch A	299
Configuring Switch B	299
Configuring Switch C	300
Verifying the RIP Configuration	301
OSPF (Open Shortest Path First)	301
OSPF Overview	301
Basic OSPF Configuration Tasks	302
Configuring OSPF Router ID	302
Configuring OSPF Areas	302
Configuring OSPF Interfaces	303
Configuring Additional OSPF Parameters	303
Basic OSPF Configuration Example	304
Configuration Example	307
Configuring Switch A	308
OSPF Stub Area NSSA Summary	310
OSPF Virtual Link Configuration Guide	310
Configure Virtual Links	310
Enable Virtual Links on the Area Border Routers	312
Checking an IPv6 OSPF	313
OSPF Area Range Configuration Guide	314
Importing an External Route into an OSPF Area	316
BFD Supports for OSPF, BGP, Static Route and ECMP	319
Configuring the Mode	319
Configuring Detect-Multiplier, Min-Receive-Interval and Min-Transmit-Interval	319
Enable BFD on L3 Interface	319
Enable BFD Supporting for OSPF4	320
Enable BFD Supporting for OSPF6	320
Enable BFD Supporting for BGP	320
Enable BFD Supporting for Static Route	320
Enable BFD Supporting for ECMP	321
BFD Configuration	321
Configuring Switch A	321
Configuring Switch B	322

Verifying the BFD Configuration	322
Configuring Switch A	323
	323
Configuring Switch B	323
Verifying the BFD Configuration	324
Configuring ECMP (Equal-Cost Multipath Routing)	324
Configuring the Equal-Cost Path Maximum	324
Configuring Static ECMP Routing	325
Configuring ECMP Hash Fields	326
Configuring VRRP (Virtual Router Redundancy Protocol)	326
Configuring Active-Active-VRRP	327
Configuring the IPv6 Neighbor Aging Time	328
Configuring a Static IPv6 Neighbor	328
Configuring IPv6 Router Advertisement	329
IPv6 Static Routing Configuration	329
Configuring a Static Route for IPv6	329
OSPFv3 Routing Protocol Configuration	330
Configuring the Router ID	330
Configuring an OSPF Area and Area-Type	330
Configuring OSPF Interfaces	330
Configuring Additional OSPF Interface Parameters	331
ACL and Filter Configuration	331
Configuring ACLs	331
Configuring ACLs in VLANs	332
Configuring ACL Discard TCP ACK	332
Configuring ACL logging for Match Statistics	333
Configuring Control Plane Security Guide	333
About Control Plane Security Policer	333
Control Plane Security Policer Implementation	334
About Rate-limit and Burst applied in firewall policer	335
Control Plane Security Rate Limits and Action Examples	336
Configuring Control Plane Security Policer	336
Configuring Control Plane Security Filter	337
Configuring Control Plane Security Applied to Inbound-Interface	338
IPv4/IPv6 BGP Configuration	338
IPv4 BGP configuration	339
BGP Configuration Guide	339
Configuring a BGP Router ID	339
Configuring BGP Local-AS	339
Configuring External BGP Peering	340
Configuring Internal BGP Peering	340

Configuring the BGP Local Preference	340
Configuring BGP MED	341
Configuring BGP Next Hop	341
Configuring BGP Route Reflectors	342
Configuring BGP Confederations	342
Configuring the BGP Connect Timer	343
Configuring MD5 Authentication for TCP Connections	343
Configuring EBGP Fast-External-Fallover	343
Configuring BGP Route Summarization	344
Configuring BGP to Advertise Default Routes to Peers	344
Configuring BGP AS Loop	345
Configuring BGP Load Balancing	345
Configuration Example	345
Configuring Switch A	346
Configuring Switch B	347
Configuring Switch C	348
Configuring Switch D	348
Viewing BGP Peer Status on Switch B	349
	350
Configuring Switch A to Advertise Route 192.168.10.0/24	350
	351
Configuring Switch B to Advertise a Connected Route	351
BGP Route Reflector Configuration Example	351
Configuring Switch A	352
Configuring Switch B	353
Configuring Switch C	354
Configuring Switch D	355
Verifying Configurations	356
BGP Confederation Configuration Example	356
Configuring Switch A	357
Configuring Switch B	358
Configuring Switch C	359
	360
Configuring Switch D	360
Configuring Switch E	361
Verifying the Configuration	361
BGP Load Balancing Configuration Example	362
Configuring Switch A	362
Configuring Switch B	363
Configuring Switch C	364
Configuring Switch D	365
Configure an EBGP Connection	365
Viewing BGP Peer Status on Switch B	366
Viewing BGP Peer Status on Switch C	366
Configuring Switch D to Advertise Route 192.168.10.0/24	367

Configuring BGP Multipath Path-Relax on Switch A	368
IPv6 BGP Configuration	368
IPv6 BGP Introduction	369
Building Peering Sessions	370
Command References	370
EBGP Peering	371
IBGP Peering	378
Establish BGP Peer Use 4-byte-AS-Number	385
Case 1	385
One Peer Enable 4-byte-as-number. Another Peer Didn't Enable 4-byte-as-number.	385
Case 2	388
Two Peer All Enable 4-byte-as-number	388
Case 3	390
Two Peer All Enable 4-byte-as-number	390
Sources of Routing Updates	392
Injecting Information Dynamically into BGP	393
Configuration Command References:	393
Configure Example 1:	393
Configure Example 2:	396
Injecting Information Statically into BGP	401
Configuration Command References:	401
BGP Attributes	404
The NEXT_HOP Attribute	404
Configuration Example 1	404
Case 1 :	404
Configuration Example 2	407
Case 2:	407
Configuration Example 3	409
Case 3:	409
The AS_PATH Attribute	410
Commands References	412
The MULTI_EXIT_DISC Attribute	416
Command References	420
Aggregation	421
Well-known Communities	421
Configuration Example 1	422
Configuration Example 2	425
Configure Example 3:	428
Command References	430
Command Reference	437
Controlling Large-Scale Autonomous System	441
Command References	442
Route Reflectors Command References	446
The Route Reflector Preserves IBGP Attributes	447

Avoiding Loops	447
Configuration Example 1:	447
Configuration Example 2:	449
Configuration Example 4:	454
BGP Load Balance Command References	457
Configure Example 1: Redundancy	458
Configuration Example 2 : Load Balance (BGP route have the same AS-path)	460
Configuration Example 3 : Load Balance (BGP route have different AS-path)	464
Route Dampening Command Reference	467
Penalty	468
Half-life	468
Suppress Limit	468
Reuse Limit	468
Max-suppress Timer	468
Configuration Example 1:	469
Configuration Example 2:	472
Label BGP	475
Labeled BGP Support	475
Restrictions	475
Overview	475
How to Configure	476
Configuration Example for Labeled Support	477
IPV4 Labeled BGP Configuration	477
Scenario 1: Incoming Three Layer Label	477
Step 1: Configure loopback interface	477
Scenario 2: Incoming Two Layer Label	483
Scenario 3: Incoming One Layer Label	489
Scenario 4: Incoming IP Packet with No-label	495
Step 1: Configure Loopback interface	495
IPV6 Labeled-BGP Configuration	501
Scenario 1:Incoming three layer label	501
Step 1:Configure Loopback interface	501
Step 2:Configure Vlan-interface	501
Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable	502
Step 4.Configure BGP	502
Step 5:Enable Labeled-BGP on switchA,B,C	503
Step 6:switchA distribute a bgp route	503
Step 7:Check BGP route table	504
Step 8:Check mpls forwarding table on switchA,B,C	504
Step 9:Configure MPLS label mapping on switchC	504
Step 10:Configure MPLS label mapping on switchA	505
Step 11:Check the mpls forwarding table on switchA ,B,C	505
Step 12:H2 send mpls packet with 3-level mpls label	505

Step13.Check the packet out of switchA, the top label 1000 will be popped at the Penultimate Hop	505
	506
Scenario 2:Incoming two layer label	506
Step 1:Configure Loopback interface	506
Step 2:Configure Vlan-interface	507
Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable	507
Step 4.Configure BGP	508
Step 5:Enable Labeled-BGP on switchA,B,C:	508
Step 6:A distribute a bgp route	509
Step 7:Check BGP route table	509
Step 8:Check mpls forwarding table on switchA,B,C:	509
Step 9:Configure MPLS label mapping on switchC	510
Step 10:Configure MPLS label mapping on switchA	510
Step 11:Check the mpls forwarding table on A ,B,C	510
Step 12:H2 send mpls packet with 2-level mpls label	510
Step 13.Check the packet out of A, the top label 1000 will be popped at the Penultimate Hop	511
Scenario 3:Incoming one layer label	511
Step 1:Configure Loopback interface	512
Step 2:Configure Vlan-interface	512
Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable	513
Step 4.Configure BGP	513
Step 5:Enable Labeled-BGP on switchA,B,C:	514
Step 6:A distribute a bgp route	514
Step 7:Check BGP route table	514
Step 8:Check mpls forwarding table on switchA,B,C:	515
Step 9:Configure MPLS label mapping on switchC	515
Step 10:Configure MPLS label mapping on switchA	515
Step 11:Check the mpls forwarding table on switchA,B,C	516
Step 12:H2 send mpls packet with 1-level mpls label	516
Step 13.Check the packet out of A, the top label 1000 will be popped at the Penultimate Hop	517
Scenario 4:Incoming IPv6 packet with no label	517
Step 1:Configure Loopback interface	517
Step 2:Configure Vlan-interface	518
Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable	518
Step 4.Configure BGP	519
Step 5:Enable Labeled-BGP on switchA,B,C:	519
Step 6:A distribute a bgp route	519
Step 7:Check BGP route table	520
Step 8:Check mpls forwarding table on switchA,B,C:	520

Step 9:Configure MPLS label mapping on switchC	521
Step 10:Configure MPLS label mapping on switchA	521
Step 11:Check the mpls forwarding table on switchA ,B,C	521
Step 12:H2 send IPv6 packet with no-label	521
Step 13.Check the packet out of switchA	522
Debugging CLI for Labeled-BGP	522
Show BGP Peers	522
Show BGP Peer Details	522
Show BGP Routes IPv4/IPv6	523
Show BGP Routes IPv4/IPv6 Detail	523
Show BGP Routes IPv4/IPv6 Label	523
Show Route Label	524
Show MPLS Forwarding-Table IPv4/IPv6 All	524
set protocols bgp trace options flag [all/ FSM/configuration/damp/event/keepalive/nexthop-resolution/notification/notify-snmp/open timer/update] disable false	524
set protocols mpls traceoptions disable false	525
Layer 3 Multicast Routing Configuration	526
IGMP Snooping Configuration	526
IGMP Snooping Querier	526
IGMP Configuration	527
Configuring an IGMP Interface	527
	527
Configuring IGMP Parameters for the IGMP Interface	528
	528
Configuring an IGMPv3 Interface	528
	528
Joining and Leaving a Group; Displaying Group Information	528
PIM-SM Configuration	529
Static RP Configuration	529
PIM-SM Configuration Example	530
Configuring Switch A	530
Configuring Switch B	531
Configuring Switch C	532
Multicast Command List	533
PIM-SSM Configuration	533
PIM-SM Basic Configuration	533
PIM-SSM Configuration Example	533
Configuring Switch A	534
Configuring Switch B	535

Configuring Switch C	535
QoS Configuration	537
Weighted Random Early Detection	537
WRED Overview	537
WRED Configuration Tasks	538
WRED Configuration Example	540
Enable WRED	540
Set Maximum and Minimum Thresholds	540
Set Drop Probability	541
Enable ECN	541
SP Queue Principle	541
WRR Principle	541
WFQ Principle	542
SP Configuration Example	542
Configure two forwarding-classes	543
Configuring classifier	543
Apply classifiers to two ingress ports	544
Generate Traffic	544
WRR Configuration Example	544
Configure Scheduler	545
Configure Two Forwarding-Classes	545
Configuring Classifiers and Apply Classifiers to Ingress Ports	546
Configuring Scheduler Profile and Apply Scheduler Profile to Egress Ports	546
Generate Traffic	547
WFQ Configuration Example	547
Configure Scheduler	548
Configure Two Forwarding-Classes	548
Configuring Classifier and Apply Classifiers to Ingress Ports	549
Configuring Scheduler Profile and Apply Classifiers to Egress Ports	550
Generate Traffic	550
QoS Configuration Guide	550
Configure a Scheduler	550
Configure a Forwarding Class	551
Configure Scheduler Profile	551
Configure Scheduler Profile to Specified Port	551
Configure a Classifier with IEEE 802.1/DSCP/ToS	552
Configure Classifier Relevant to Forwarding Class	552
Configure Classifier to Specified Port	552
PFC Configuration Example	553
Configure PFC Profile	553

Apply PFC Profile to Port	553
Show PFC Frame Statistics on Port	553
OpenFlow in Crossflow Mode	556
Configure ovs-vswitchd Locally	556
Configure ovssdb-server Remotely	556
Configure ovs-vswitchd by Remote Server 10.10.50.42	557
Crossflow Mode Introduction	557
CrossFlow Mode Introduction	558
Enable Crossflow mode and local-control disable (Openflow port):	559
Enable C rossflow mode and local-control enable (Crossflow port):	559
Disable C rossflow mode (Legacy port):	559
Install a flow with action "NORMAL":	561
CrossFlow Mode Known Limitations	561
Conflict Management between OpenFlow and CLI Configurations	561
Default Drop In OpenFlow	562
Other Limitations	562
Configurations Example in Crossflow Mode	562
Configure the OpenFlow Port in CrossFlow Mode	562
Configure the Hybrid Port in CrossFlow Mode	563
Examples	563
Basic Configurations	563
Lag Configurations	564
Flow Priority Configurations	565
FDB Configurations	566
Route Configurations	567
VXLAN Configuration	569
About VXLAN	569
VXLAN Technology	569
VXLAN Standards	569
VXLAN Inner 802.1Q	570
Encapsulation mode	570
Decapsulation-mode	570
VXLAN ECMP	572
VXLAN Mac Learning	572
VXLAN Traffic Forwarding	573
Unicast Traffic	573
Broadcast and Unknown T raffic	573
Configure map port to VXLAN VNI or map port&vlan to VXLAN VNI Step	573

Configure map vlan to VXLAN VNI Step	574
Topology	575
SW1 Configure	575
SW2 Configure	575
SW3 Configure	576
VXLAN ECMP Configuration	576
Requirements	576
Overview	576
Topology	577
Configuring VXLAN on Trident-II Series Switches	577
CLI Quick Configuration	577
SW1 Configure	577
SW2 Configure	578
Configuring VXLAN Step-by-Step Procedure	578
Configure VXLAN Step-by-Step for SW2	579
View the VXLAN Table of SW1	580
OVSDB VTEP Configuration	581
Configuring an OVSDB VTEP	581
Overview	581
Configuring an OVSDB Connection with a Controller	581
Creating and Installing an SSL Key and Certificate on a Pica8 Device	581
Configure a OVSDB VTEP Step	582
OVSDB VTEP with Midonet Configuration	583
Requirements	583
Overview	583
Topology	583
Configuring OVSDB on Trident-II Series Switches	584
CLI Quick Configuration	584
SW1 Configure :	584
Configuring OVSDB Step-by-Step Procedure	585
Configuring OVSDB Step-by-Step for SW1	585
Manually Configuring the Switch as a VTEP Gateway by Midonet cli	586
Configuring the ovbdb by midonet cli for SW1	586
View the VXLAN table of SW1:	588
View the VXLAN mac table of SW1:	588
Dump the ovbdb hardware vtep table of SW1:	588
OVSDB VTEP with NSX Configuration	590
Requirements	590
Overview	591
Topology	591
Configuring OVSDB on Trident-II Series Switches	591
CLI Quick Configuration	591
SW1 Configure	591
Configuring OVSDB Step-by-Step Procedure	592

Configure OVSDB Step-by-Step for SW1	592
Manually configuring the Switch as a VTEP Gateway by NSX Manager	593
OVSDB VTEP with vtep-ctl Configuration Examples	598
Requirements	598
Overview	598
Topology	598
Configuring OVSDB on Trident-II Series switches	599
CLI Quick Configuration	599
SW1 Configure:	599
SW2 Configure:	599
Configuring OVSDB Step-by-Step Procedure	599
Configure OVSDB Step-by-Step for SW1	600
Configure OVSDB Step-by-Step for SW2	600
Manual configuration of the Switch as a VTEP Gateway by vtep-ctl	601
Configuring the ovbdb by vtep-ctl for SW1	601
Configuring the ovbdb by vtep-ctl for SW2	601
View the VXLAN table of SW1:	601
View the VXLAN mac table of SW1:	602
L2/L3 Troubleshooting Guide	603
Monitoring and Debugging L2/L3 protocols	603
Find and Configure the Log File	603
Enable Important Debugs	603
Find the Core Dump File	605
	606
Find the last_death file for Troubleshooting	606
Routing and Forwarding Table	606
Pipe () Filter Functions	606
Comparing Configurations	607
Counting the Number of Output Lines	607
Displaying Output that Matches a Pattern	607
Omitting Output that Matches a Pattern	608
Preventing Output from Being Paginated	608
Show Tech-Support Command	608
Configuration Appendix	610
Other Command List	610

Pica8 PicOS supports Layer 2 switching protocols, including: STP, RSTP, MSTP, MAC learning, and Q-in-Q. PicOS also supports several Layer 3 protocols, including: static routing, RIPv2, OSPF, IGMP, PIM-SM, and IPv6. This guide provides instructions and examples for configuring switches and controllers. Intended for system administrators, this guide assumes a working knowledge of Layer 2 and Layer 3 protocols.

PicOS can run in 2 different modes of operation:

- **Open vSwitch (OVS) Mode:** In this mode, PicOS is dedicated and optimized for OpenFlow applications.
- **Layer 2/Layer 3 (L2/L3) Mode:** In this mode, PicOS can run switching and routing protocols, as well as OpenFlow applications.

In OVS mode, L2/L3 daemons are not running, and the system is fully dedicated to OpenFlow and OVS. In L2/L3 mode, L2/L3 daemons are running, and OVS can also be used if CrossFlow is activated.

This chapter assumes that user is running PicOS L2/L3 mode. Please see PicOS Mode Selection to learn how to switch between L2/L3 and OVS modes.

Supported Features

Command-Line Interface

- From Linux Shell to L2/L3 Shell
- Operation Mode and Configuration Mode
- Displaying the Current Configuration
- Display Setting Configuration
- Rolling Back a Configuration
- Managing Configuration Files
- Saving and Loading Configuration Files
- Commit Confirmed
- Commit Check
- Commit Failed and Exit Discard
- Configuring a Command Alias
- Configure L2/L3 from Linux Shell
- Bash Linux Shell
- PicOS Upgrade and Configuration Change
- Set CLI

System Administration

- Configuring DHCP and Static IP Address
- Configuring DHCP Relay
- Configuring DHCP Relay after PicOS 2.9.2(including 2.9.2)
- Configuring DHCP Option82
- Configuring DHCP Snooping
- Configuring DHCP Snooping after PicOS 2.9.2(including 2.9.2)
- Configuring DHCP Relay and DHCP Snooping together after PicOS 2.9.2(including 2.9.2)
- Configuring a User Account
- Configuring Authentication, Authorization, and Accounting
- Configuring SSH and Telnet Parameters
- Configuring the Log-in ACL
- Configuring NTP and the Time Zone Parameter
- Configuring PTP
- Configuring the linux-config-unreliable
- Configuring IPFIX
- Configuring sFlow
- Configuring SNMPv2
- Configuring the Syslog Log Level
- Configuring the Syslog Disk and Syslog Host
- Displaying System Information
- Technical Support
- Flushing ARP and the Neighbor Table
- Rebooting the System
- Displaying the Debugging Message
- System Management Command List
- IPv6 Management Support
- PoE Configuration
- NETCONF Configuration
- Configuring SNMPv3

Ethernet Ports Configuration

- Physical Ethernet Port Configuration
- VLAN Port Configuration
- 40G Changes to 4*10G in L2/L3
 - 40G Changes to 4*10G in L2/L3 mode on P-5101
 - 40G Changes to 4*10G in L2/L3 mode on P-5401
 - 40G Changes to 4*10G in L2/L3 mode on accton_as6701_32x
 - 40G Changes to 4*10G in L2/L3 mode on P-3922
 - 40G Changes to 4*10G in L2/L3 mode on P-3920
 - 40G Changes to 4*10G in L2/L3 mode on accton_as5712_54x/HP5712
 - 40G Changes to 4*10G in L2/L3 mode on P-3930
 - 40G Changes to 4*10G in L2/L3 mode on Niagara2632XL
 - 40G Changes to 4*10G in L2/L3 mode on Niagara2948_6XL
 - 40G Changes to 4*10G in L2/L3 mode on as5812_54t
 - 40G Changes to 4*10G in L2/L3 mode on arctica4806xp
- LLDP Configuration (Link Layer Discovery Protocol)
- Voice Vlan with LLDP Configuration Guide
- Voice Vlan with LLDP Compliance CDP Configuration Guide
- Voice Vlan Configuration Guide
- LLDP MED Configuration
- Voice Vlan with LLDP MED Configuration Guide
- Static Link Aggregation (LAG) Configuration
- LAG Hashing Configuration
 - LAG Hashing Configuration and Example
 - LAG Hash Mapping
 - Resilient LAG Hashing Configuration and Example
- Link Aggregation Control Protocol (LACP) Configuration
- Storm Control in Ethernet Port Configuration
- UDLD Configuration
- Port Security Configuration
- 802.1x Configuration
- Buffer management

Layer 2 Switching Configuration

- Static MAC entries and Dynamic MAC Address Learning
- Cut-Through Switching Method
- MLAG Configuration Guide
 - Configuring a Basic MLAG Step-by-Step Procedure
 - Configuring a Basic MLAG Example
 - Configuring Switch A with Static and LACP LAG
 - Configuring Switch B with Static and LACP LAG
 - Configuring Switch C with LACP and LAG
 - View the MLAG Internal and Neighbor Status
 - Configuring a MLAG domain with MSTP example
 - Configuring Switch A with LACP LAG
 - Configuring Switch B with LACP LAG
 - Configuring an Aggregation Interface to VLAN Members
 - Configuring Switch C and Switch D with LACP LAG
 - Configure Mlag and Active-Active-Vrrp example
 - Configure on Switch A
 - Configure on Switch B
 - Configure on Switch C
 - Configure on Switch D
 - Configuring a MLAG with vxlan Example
 - 01 Configuring mlag in Switch A and Switch B
 - Configuring Lag in Switch C and Switch D
 - 02 Configuring Lag in Switch C and Switch D
 - 03 Configuring vxlan in Switch A and Switch B
- Q-in-Q Basic Port Configuration
- MSTP Configuration
- MSTP Configuration Example
- PVST Configuration
- PVST Configuration Example
- Configuring Mirroring
- Configuring Mirroring Guide

- BPDU Tunneling Configuration
- Configuring IPv6 RA Guard
- LFS Configuration
- Symmetric Hash for ECMP and LAG Configuration Example

Layer 3 Unicast Routing Configuration

- Layer 3 VLAN Interface Configuration
- ARP Configuration
- Dynamic ARP Inspection
- Default Administrative Distance Values
- Static Routing Configuration
- Static Routing Configuration Example
- RIPv2 Routing Protocol Configuration
- RIPv2 Routing Configuration Example
- OSPF (Open Shortest Path First)
 - OSPF Overview
 - Basic OSPF Configuration Tasks
 - Basic OSPF Configuration Example
- OSPF Configuration Example: NSSA Stub Normal
- OSPF Stub Area NSSA Summary
- OSPF Virtual Link Configuration Guide
- OSPF Area Range Configuration Guide
- Importing an External Route into an OSPF Area
- BFD Protocol Configuration
- BFD Basic Configuration Example
- Configuring ECMP (Equal-Cost Multipath Routing)
- Configuring VRRP (Virtual Router Redundancy Protocol)
- Configuring Active-Active-VRRP(Active Active Virtual Router Redundancy Protocol)
- IPv6 Neighbor Configuration
- IPv6 Static Routing Configuration
- OSPFv3 Routing Protocol Configuration
- ACL and Filter Configuration
 - Configuring Control Plane Security Guide
 - Configuration Control Plane Security Policer
 - Control Plane Security Examples
- IPv4/IPv6 BGP Configuration

- IPv4 BGP configuration
 - BGP Configuration Guide
 - BGP Basic Configuration Example
 - BGP Route Reflector Configuration Example
 - BGP Confederation Configuration Example
 - BGP Load Balancing Configuration Example
- IPv6 BGP Configuration
 - IPv6 BGP Introduction
 - Building Peering Sessions
 - EBGP Peering
 - IBGP Peering
 - Establish BGP Peer Use 4-byte-AS-Number
 - Sources of Routing Updates
 - Injecting Information Dynamically into BGP
 - Injecting Information Statically into BGP
 - BGP Attributes
 - The NEXT_HOP Attribute
 - The AS_PATH Attribute
 - The LOCAL_PREF Attribute
 - The MULTI_EXIT_DISC Attribute
 - The COMMUNITY Attribute
 - BGP-4 Aggregation
 - Synchronization
 - Controlling Large-Scale Autonomous System
 - Confederations
 - Route Reflectors
 - Redundancy and Load Balancing
 - Designing Stable Internets
- Label BGP
 - Labeled BGP Support
 - Configuration Example for Labeled Support
 - IPV4 Labeled BGP Configuration
 - IPV6 Labeled-BGP Configuration
 - Debugging CLI for Labeled-BGP

Layer 3 Multicast Routing Configuration

- IGMP Snooping Configuration
- IGMP Configuration
- PIM-SM Configuration
- PIM-SM Configuration Example
- Multicast Command List
- PIM-SSM Configuration
- PIM-SSM Configuration Example

QoS Configuration

- Weighted Random Early Detection
 - WRED Overview
 - WRED Configuration Tasks
 - WRED Configuration Example
- QoS Principle
- SP Configuration Example
- WRR Configuration Example
- WFQ Configuration Example
- QoS Configuration Guide
- PFC Configuration Example

OpenFlow in Crossflow Mode

- Crossflow Mode Introduction
- CrossFlow Mode Known Limitations
- Configurations Example in Crossflow Mode

VXLAN Configuration

- [VXLAN Configuration Guide](#)
- [VXLAN Base Configuration Example](#)
- [VXLAN ECMP Configuration](#)
- [OVSDB VTEP Configuration](#)
 - [Configuring an OVSDB VTEP](#)
 - [OVSDB VTEP with Midonet Configuration](#)
 - [OVSDB VTEP with NSX Configuration](#)
 - [OVSDB VTEP with vtep-ctl Configuration Examples](#)

L2/L3 Troubleshooting Guide

- Monitoring and Debugging L2/L3 protocols
- Routing and Forwarding Table
- Using Pipe (|) Filter Functions
- Using the show tech-support Command

Configuration Appendix

- Other Command List

The PicOS documentation is available at the Pica8 website:

<http://www.pica8.com/portal/>

Summary of Supported Services

Table 1 Supported Layer 2 and Layer 3 Protocols

Category	Features
System Management and Administration	Support for clock/date setting and NTP (Network Time Protocol)
	Support for inbound IP access via any routed interface
	Support for DHCP (Dynamic Host Configuration Protocol); DHCP client, DHCP relay, DHCP Option82, and DHCP snooping
	Support for multiple local user accounts
	Support for SSHv2 (Secure Shell) protocol
	Ability to enable debugging for a specific module
	Support for Read Only and Read Write access SNMP (Simple Network Management Protocol)
	Support for IPFIX (IP Flow Information Export), monitors data flow in specified server
Device Configuration, Software, and File Management	Support the ability to save the configuration to flash on the device
	Support for configuration versioning and rollback; compares the two configurations, identifying differences
	Ability to import/export configuration files, device software, and logs from a file on a remote server (tftp/scp as options)
	Ping and Trace route tool from CLI (command line interface)
	SSH and telnet tool from CLI
	Ability to view and configure MAC/ARP (Address Resolution Protocol) table information
Layer 2 Forwarding and Protocol	Support for LLDP (Link Layer Discovery) protocols for detecting devices on a link
	Support for LACP (Link Aggregation Control) protocol and hashing of traffic using src/Dst (Source/Destination) MAC address, Src/Dst IP address, and Layer 4 port information and flag
	Support for 802.1q trunked interfaces, for both single and LAG (Link Aggregation Group) interfaces
	Support for 802.1q tagged/untagged interfaces and native tags
	Support for Q-in-Q
	Support for Jumbo Frame
	Support for 802.1d STP (Spanning Tree Protocol)
	Support for 802.1w RSTP (Rapid STP) and PVST (Per-VLAN STP)
	Support for 802.1s MSTP (Multiple Spanning Tree protocol)

	Support for functionality of BPDU (Bridge Protocol Data Unit) Guard / Filter/UDLD (Unidirectional Link Detection)
	Support for storm-control for unicast, multicast, broadcast
	Support for ingress/egress port mirroring
	Support for 802.1p in Layer 2 forwarding
	Support for Flow control per-interface
	Support for IGMP (Internet Group Management Protocol) snooping enable per-VLAN
	Support for IGMP snooping query per-VLAN
Layer 3 Forwarding and Routing Protocol	Full support for dual stacked IPv4 and IPv6 addressing.
	Support for 6 members in a Layer 3 LAG (Link Aggregation Group) interface
	Support for IPv4 and IPv6 static route configuration
	Support for OSPFv2 (Open Shortest Path First) IPv4 only
	Support for stub, normal, and NSSA (Not-So-Stubby Area) OSPF area types
	Support for up to 32 equal-cost routes in OSPF
	Support for RIP routing protocol
	Support for BGP (Border Gate Protocol) routing and BFD (Bidirectional Forwarding Detection)
	Support for 128 equal-cost routes in the device's routing/forwarding tables
	Support for ECMP (Equal-Cost Multi-path) routing with hashing of traffic using Src/Dst IP and Port
	Support the ToS and DSCP (Differentiated Services Code Point) in Layer 3 forwarding
	Support for IGMP v1/v2
	Support for PIM-SM (Protocol Independent Multicast Routing-Sparse Mode)
	Support for VRRP (Virtual Router Redundancy Protocol)

Command-Line Interface

CLI Configuration

This chapter describes the different ways to configure PicOS and demonstrates the CLI configuration.

There are 2 CLI's used to configure PicOS:

- The Linux CLI
- The PicOS CLI

The Linux CLI is a standard debian based bash shell.

A good Bash tutorial can be found at this address:

<http://www.tldp.org/LDP/Bash-Beginners-Guide/html/>

PicOS added some commands to the standard Bash shell:

Version - This is to provide the PicOS version running on the switch

```
admin@XorPlus$version
Copyright (C) 2009-2014 Pica8, Inc.
=====
Hardware Model           : P-5101
Linux System Version/Revision : 2.5/17907
Linux System Released Date  : 10/14/2014
L2/L3 Version/Revision     : 2.5/17907
L2/L3 Released Date        : 10/14/2014
OVS/OF Version/Revision    : 2.5/17907
OVS/OF Released Date       : 10/14/2014
```

cli - command to move to the PicOS command or launch PicOS CLI commands from the Linux shell.

```
admin@XorPlus$cli
Synchronizing configuration...OK.
Pica8 PicOS Version 2.5
Welcome to PicOS L2/L3 on XorPlus
admin@XorPlus>
```

In the above example, the cli command is used to move to the PicOS CLI.

```
admin@XorPlus$cli -c "show version"
Synchronizing configuration...OK.
Pica8 PicOS Version 2.5
Welcome to PicOS L2/L3 on XorPlus
admin@XorPlus>
Execute command: show version
.
Copyright (C) 2009-2014 Pica8, Inc.
=====
Base ethernet MAC Address   : 48:6e:73:01:00:01
Hardware Model              : P-5101
Linux System Version/Revision : 2.5/17907
Linux System Released Date  : 10/14/2014
L2/L3 Version/Revision     : 2.5/17907
L2/L3 Released Date        : 10/14/2014
```

In the command above, the `cli` command is used to launch commands of the PicOS CLI from the Linux shell.

- From Linux Shell to L2/L3 Shell
- Operation Mode and Configuration Mode
- Displaying the Current Configuration
- Display Setting Configuration
- Rolling Back a Configuration
- Managing Configuration Files
- Saving and Loading Configuration Files
- Commit Confirmed
- Commit Check
- Commit Failed and Exit Discard
- Configuring a Command Alias
- Configure L2/L3 from Linux Shell
- Bash Linux Shell
- PicOS Upgrade and Configuration Change
- Set CLI

From Linux Shell to L2/L3 Shell

Once in the Linux shell, user can use the command `pica_sh` or `cli` (under `/pica/bin`) to launch the L2/L3 CLI (or XORP CLI).

```
admin@Lima$
admin@XorPlus$cli
Synchronizing configuration...OK.
Pica8 PicOS Version 2.4
Welcome to PicOS L2/L3 on XorPlus
admin@XorPlus>
```

To come back to the Linux Shell from the L2/L3 CLI (or XORP CLI), use the `exit` command.

```
admin@XorPlus> exit
admin@XorPlus$
```

Operation Mode and Configuration Mode

By default, the switch's operation mode is activated when it starts up.

Welcome to PicOS L2/L3 on XorPlus

```
admin@XorPlus>
```

Activate the configuration mode by entering the `configure` command. Be sure to enter the configuration mode if the `admin@XorPlus#` prompt appears.

```
admin@XorPlus> configure
Entering configuration mode.
There are no other users in configuration mode.
admin@XorPlus#
```

Displaying the Current Configuration

In L2/L3, non-default configuration can be displayed with the **show** commands. The command **show all** displays the default value of the current configuration. Default configurations are shown in the `pica_default.boot` file. The command **show running-config** displays the configuration active on the system.

```
admin@XorPlus# show
vlans {
vlan-id 200 {
}
}
admin@XorPlus#
admin@XorPlus# show all
vlans {
vlan-id 200 {
description: ""
vlan-name: "default"
l3-interface: ""
}
}
```

```
admin@XorPlus> show running-config
  vlans {
    vlan-id 200 {
    }
  }
```

Display Setting Configuration

This command displays which settings the configuration has set by default and which settings need to set manually.

```
admin@XorPlus# show | display set
set interface ethernet-switching-options analyzer test input ingress ge-1/1/2
set interface ethernet-switching-options analyzer test input egress ge-1/1/2
set interface ethernet-switching-options analyzer test output "ge-1/1/3"
```

```

admin@XorPlus# show | display set
    set vlans vlan-id 11
    set vlans vlan-id 22
admin@XorPlus#
admin@XorPlus# set vlans vlan-id 33
admin@XorPlus# set vlans vlan-id 44
admin@XorPlus# set vlans vlan-id 55
admin@XorPlus# show | display set
    set vlans vlan-id 11
    set vlans vlan-id 22
>   set vlans vlan-id 33
>   set vlans vlan-id 44
>   set vlans vlan-id 55
admin@XorPlus# com
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# show | display set
    set vlans vlan-id 11
    set vlans vlan-id 22
    set vlans vlan-id 33
    set vlans vlan-id 44
    set vlans vlan-id 55
admin@XorPlus#

```

Rolling Back a Configuration

Each time a configuration in L2/L3 is committed, a rollback configuration file is created. For example, if configuration is committed 10 times, **pica.conf.01** through **pica.conf.10** are created. User can rollback to any of these configurations when necessary. The maximum quantity of rollback files is limited to 50. The current configuration is located in **pica.conf**.

```

admin@XorPlus# rollback 1
admin@XorPlus# Loading config file...
Config file was loaded successfully.
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Displaying the difference between the current config with destination config file

```

admin@XorPlus# show | compare rollback 2
[edit vlans]
-----
+vlan-id 3 {
+}
admin@XorPlus#

```

Managing Configuration Files

Configuration files can be copied, deleted, or renamed in the system, but system files should *not* be deleted.

The L2/L3 configuration is stored in /pica/config/pica_startup.boot.

```

admin@XorPlus$
admin@XorPlus$cd /pica/config/
admin@XorPlus$ls
admin          pica.conf.06  pica.conf.13  pica.conf.20  pica.conf.27  pica.conf.34
pica.conf.41  pica.conf.48
pica.conf     pica.conf.07  pica.conf.14  pica.conf.21  pica.conf.28  pica.conf.35
pica.conf.42  pica.conf.49
pica.conf.01  pica.conf.08  pica.conf.15  pica.conf.22  pica.conf.29  pica.conf.36
pica.conf.43  pica_startup.boot
pica.conf.02  pica.conf.09  pica.conf.16  pica.conf.23  pica.conf.30  pica.conf.37
pica.conf.44
pica.conf.03  pica.conf.10  pica.conf.17  pica.conf.24  pica.conf.31  pica.conf.38
pica.conf.45
pica.conf.04  pica.conf.11  pica.conf.18  pica.conf.25  pica.conf.32  pica.conf.39
pica.conf.46
pica.conf.05  pica.conf.12  pica.conf.19  pica.conf.26  pica.conf.33  pica.conf.40
pica.conf.47

```

User can display the files of a specified directory:

```

admin@XorPlus> file list /
drwxr-xr-x 2 root xorp 4096 Sep 25 00:54 bin
drwxr-xr-x 2 root xorp 4096 Sep 24 06:21 boot
drwxr-xr-x 2 root xorp 4096 Sep 23 17:05 cftmp
-rwxr-xr-x 1 root xorp 40559 Sep 23 17:05 config.bcm
drwxr-xr-x 4 root root 4096 Sep 25 00:54 dev
drwxr-xr-x 7 root xorp 4096 Sep 25 00:55 etc
drwxr-xr-x 4 root xorp 4096 Sep 24 06:21 lib
lrwxrwxrwx 1 root root 11 Sep 24 06:21 linuxrc -> bin/busybox
drwxr-xr-x 5 root xorp 4096 Sep 24 06:21 mnt
drwxr-xr-x 2 root xorp 4096 Sep 23 17:05 opt
drwxr-xr-x 5 root xorp 4096 Sep 24 06:21 ovs
drwxr-xr-x 14 root xorp 4096 Sep 24 06:23 pica
dr-xr-xr-x 52 root root 0 Jan 1 1970 proc
-rwxr-xr-x 1 root xorp 59012 Sep 23 17:05 rc.soc
drwxr-xr-x 2 root xorp 4096 Sep 24 06:21 sbin
drwxr-xr-x 11 root root 0 Jan 1 1970 sys
drwxrwxrwx 8 root xorp 1024 Sep 25 00:55 tmp
drwxr-xr-x 7 root xorp 4096 Sep 24 06:22 usr
drwxr-xr-x 7 root xorp 4096 Sep 24 06:23 var
admin@XorPlus> file list /tmp
drwxrwxr-x 5 root xorp 1024 Sep 25 00:54 home
drwxrwxr-x 2 root xorp 1024 Sep 25 00:54 log
drwx----- 2 root root 12288 Sep 25 00:54 lost+found
drwxrwxr-x 3 root xorp 1024 Sep 25 00:55 run
drwxrwxr-x 2 root xorp 1024 Sep 25 00:54 snmp
drwxrwxr-x 2 root xorp 1024 Sep 25 00:56 system

```

Display the contents of a specified file:

```
-- 1 root root 410 Sep 24 06:23 boot.lst
-rw-rw-r-- 1 root xorp 16006 Sep 24 07:44 pica.conf
-rw-rw-r-- 1 root xorp 16003 Sep 24 07:22 pica.conf.01
-rw-rw-r-- 1 root xorp 15826 Sep 24 07:19 pica.conf.02
-rw-rw-r-- 1 root xorp 15536 Sep 24 07:18 pica.conf.03
-rw-rw-r-- 1 root xorp 15915 Sep 24 07:18 pica.conf.04
-rw-rw-r-- 1 root xorp 15567 Sep 24 07:09 pica.conf.05
-rw-rw-r-- 1 root xorp 15188 Sep 24 06:44 pica.conf.06
-rw-rw-r-- 1 root xorp 14953 Sep 24 06:35 pica.conf.07
drwxrwxrwx 2 root root 4096 Sep 24 06:25 root
admin@XorPlus> file show /pica/config/pica.conf
/*XORP Configuration File, v1.0*/
interface {
  ecmp {
    max-path: 4
    hash-mapping {
      field {
        ingress-interface {
          disable: false
        }
      }
      vlan {
        disable: false
      }
      ip-protocol {
        disable: false
      }
      ip-source {
        disable: false
      }
      ip-destination {
        disable: false
      }
      port-source {
        disable: false
      }
      port-destination {
        disable: false
      }
    }
  }
}
```

User can also copy, archive, checksum, compare, rename, and sync files.

```

admin@XorPlus> file list /pica/config
-rw-r--r-- 1 root root 410 Sep 24 06:23 boot.lst
-rw-rw-r-- 1 root xorp 16006 Sep 24 07:44 pica.conf
-rw-rw-r-- 1 root xorp 16003 Sep 24 07:22 pica.conf.01
-rw-rw-r-- 1 root xorp 15826 Sep 24 07:19 pica.conf.02
-rw-rw-r-- 1 root xorp 15536 Sep 24 07:18 pica.conf.03
-rw-rw-r-- 1 root xorp 15915 Sep 24 07:18 pica.conf.04
-rw-rw-r-- 1 root xorp 15567 Sep 24 07:09 pica.conf.05
-rw-rw-r-- 1 root xorp 15188 Sep 24 06:44 pica.conf.06
-rw-rw-r-- 1 root xorp 14953 Sep 24 06:35 pica.conf.07
drwxrwxrwx 2 root root 4096 Sep 24 06:25 root
admin@XorPlus> file copy /pica/config/pica.conf
Possible completions:
<destination-file> Copy files to and from the router
admin@XorPlus> file copy /pica/config/pica.conf /pica/config/ychen.conf
admin@XorPlus> file list /pica/config
-rw-r--r-- 1 root root 410 Sep 24 06:23 boot.lst
-rw-rw-r-- 1 root xorp 16006 Sep 24 07:44 pica.conf
-rw-rw-r-- 1 root xorp 16003 Sep 24 07:22 pica.conf.01
-rw-rw-r-- 1 root xorp 15826 Sep 24 07:19 pica.conf.02
-rw-rw-r-- 1 root xorp 15536 Sep 24 07:18 pica.conf.03
-rw-rw-r-- 1 root xorp 15915 Sep 24 07:18 pica.conf.04
-rw-rw-r-- 1 root xorp 15567 Sep 24 07:09 pica.conf.05
-rw-rw-r-- 1 root xorp 15188 Sep 24 06:44 pica.conf.06
-rw-rw-r-- 1 root xorp 14953 Sep 24 06:35 pica.conf.07
drwxrwxrwx 2 root root 4096 Sep 24 06:25 root
-rw-rw-r-- 1 root root 16006 Sep 25 02:22 ychen.conf
admin@XorPlus>
admin@XorPlus> file rename /pica/config/ychen.conf /pica/config/ychen-1.conf
admin@XorPlus> file list /pica/config
-rw-r--r-- 1 root root 410 Sep 24 06:23 boot.lst
-rw-rw-r-- 1 root xorp 16006 Sep 24 07:44 pica.conf
-rw-rw-r-- 1 root xorp 16003 Sep 24 07:22 pica.conf.01
-rw-rw-r-- 1 root xorp 15826 Sep 24 07:19 pica.conf.02
-rw-rw-r-- 1 root xorp 15536 Sep 24 07:18 pica.conf.03
-rw-rw-r-- 1 root xorp 15915 Sep 24 07:18 pica.conf.04
-rw-rw-r-- 1 root xorp 15567 Sep 24 07:09 pica.conf.05
-rw-rw-r-- 1 root xorp 15188 Sep 24 06:44 pica.conf.06
-rw-rw-r-- 1 root xorp 14953 Sep 24 06:35 pica.conf.07
drwxrwxrwx 2 root root 4096 Sep 24 06:25 root
-rw-rw-r-- 1 root root 16006 Sep 25 02:22 ychen-1.conf
admin@XorPlus>
admin@XorPlus> file checksum /pica/config/ychen-1.conf
3559192236 16006 /pica/config/ychen-1.conf
admin@XorPlus>
admin@XorPlus> file sync
admin@XorPlus>
admin@XorPlus> file compare /pica/config/pica.conf /pica/config/pica.conf.01
admin@XorPlus> file compare /pica/config/pica.conf /pica/config/pica.conf.01
3c3
< /*Last commit      : Mon Jan 13 14:13:01 2014 by admin*/
---
> /*Last commit      : Mon Jan 13 14:12:26 2014 by admin*/
510,514d509
< controller 1 {
< protocol: "tcp"
< address: 10.10.50.47
< port: 6633
< }

```

User can change the current directory using functions like `cwd` or `cd`.

```

admin@XorPlus> file cwd
Current working directory: /tmp/home/admin
admin@XorPlus>
admin@XorPlus> file cwd /pica/config
admin@XorPlus> file cwd
Current working directory: /pica/config
admin@XorPlus>

```

Saving and Loading Configuration Files

The Pica8 PicOS provides several commands to save, load, and execute PicOS configuration files, as detailed in this document.

 The **save** command saves configuration files in the `/pica/config/admin` directory. The **load** and **execute** commands look for configuration files in the same directory.

You can use the **save** command in L2/L3 configuration mode to save the running configuration to a file. The following example demonstrates how to save the running configuration to a file named *myconfig.conf*.

```

admin@Spine1# save ?
Possible completions:
 <file name>                Save running configuration to specified file
admin@Spine1# save myconfig.conf
Save done.

```

The configuration saved earlier in a file can be loaded or applied using the **load merge** or **load override** command in L2/L3 mode.

```

admin@Spine1# load ?
Possible completions:
 merge                        Merge the loaded configuration to the current running
configuration
 override                     Override the current running configuration with the loaded
configuration

```

The **load override** command completely replaces the running configuration with the configuration in a file (saved earlier). The following example replaces the running configuration with the configuration in a file named *myconfig.conf*.

```

admin@Spine1# load override ?
Possible completions:
 <text>                        Local file name
 myconfig.conf                 Size: 16643, Last changed: Wed Sep 9 21:33:21 2015
admin@Spine1# load override myconfig.conf
Loading config file...
Config file was loaded successfully.

```

The **load merge** command merges the configuration in a file (saved earlier) with the running configuration. The following example merges the configuration in a file named *myconfig.conf* with the running configuration.

```
admin@Spine1# load merge ?
Possible completions:
  <text>                               Local file name
  myconfig.conf                         Size: 16643, Last changed: Wed Sep  9 21:33:21 2015
admin@Spine1# load merge myconfig.conf
Loading config file...
Config file was applied successfully.
```

Commit Confirmed

User can commit a candidate configuration before this configuration becomes permanent. By using **commit confirmed**, the system will apply the configuration with a ten minute default. After ten minutes, the system will roll back to the original configuration automatically. User can configure the roll back time in the CLI. By default, roll back time is 10 minutes.

By default, it will be automatically rolled back to the previous configuration after 600 seconds (10 minutes).

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# commit confirmed
Merging the configuration.
Will be automatically rolled back in 600 seconds unless confirmed by new commit.
Commit OK.
admin@XorPlus#
```

Modify the rollback confirmation time

```
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# commit confirmed 100
Merging the configuration.
Will be automatically rolled back in 100 seconds unless confirmed by new commit.
Commit OK.
admin@XorPlus#
```

Commit Check

User can check to see if the configuration is correct before this configuration becomes permanent. Use **commit check**. A result of "Commit check ok" means the configuration is correct as is. Then, user can commit the configuration to become permanent. A result of "Commit check failed" means the configuration has an error that needs to be corrected.

Commit Check one correct configuration

```

admin@XorPlus# set vlans vlan-id 4094
admin@XorPlus# commit check
Commit check OK.
admin@XorPlus#
admin@XorPlus# show all
    vlans {
>     vlan-id 4094 {
>         description: ""
>         vlan-name: "default"
>         l3-interface: ""
>     }
}
admin@XorPlus#
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# show all
    vlans {
        vlan-id 4094 {
            description: ""
            vlan-name: "default"
            l3-interface: ""
        }
    }
admin@XorPlus# run show vlans
VlanID   Tag           Interfaces
-----
1        untagged      ge-1/1/1, ge-1/1/2, ge-1/1/3, ge-1/1/4, ge-1/1/5
          ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9, ge-1/1/10
          ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14, ge-1/1/15
          ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19, ge-1/1/20
          ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24, ge-1/1/25
          ge-1/1/26, ge-1/1/27, ge-1/1/28, ge-1/1/29, ge-1/1/30
          ge-1/1/31, ge-1/1/32, ge-1/1/33, ge-1/1/34, ge-1/1/35
          ge-1/1/36, ge-1/1/37, ge-1/1/38, ge-1/1/39, ge-1/1/40
          ge-1/1/41, ge-1/1/42, ge-1/1/43, ge-1/1/44, ge-1/1/45
          ge-1/1/46, ge-1/1/47, ge-1/1/48, ge-1/1/49, ge-1/1/50
          te-1/1/51, te-1/1/52
4094     tagged
          untagged
          tagged

```

Commit Check one incorrect configuration

```

admin@XorPlus# set vlans vlan-id 4096
admin@XorPlus# commit check
Invalid vlan 4096
Commit check failed.
admin@XorPlus#
admin@XorPlus# show
    vlans {
>     vlan-id 4096 {
>     }
}
admin@XorPlus# commit
Invalid vlan 4096
Commit failed.
admin@XorPlus#
admin@XorPlus# exit discard
admin@XorPlus>
admin@XorPlus> configure

```

Commit Failed and Exit Discard

Switch from the configuration mode to the execution mode *without* any uncommitted configurations.

```
admin@XorPlus# exit
admin@XorPlus>
```

Use the exit discard command to enter the execution mode from the configuration mode with any uncommitted or failed committed configurations.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable true
admin@XorPlus# exit
ERROR: There are uncommitted changes.
Use "commit" to commit the changes, or "exit discard" to discard them.
admin@XorPlus# exit discard
admin@XorPlus>
```

Configuring a Command Alias

User can configure an alias for a PicOS command. This CLI also supports multiple parameters. In other words, user can use the parameter in PicOS CLI e.g. \$1,\$2....., which will be used in the alias command.

```
admin@XorPlus# set alias set_vlans as "set vlans vlan-id $1"
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set_vlans 10
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set alias set_vlans_interface as " set vlans vlan-id $1 vlan-name $2"
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set_vlans_interface 20 vlan20
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure L2/L3 from Linux Shell

Configure L2/L3 via Linux Shell on PicOS.

```

admin@XorPlus$cli -c "configure;set vlans vlan-id 100;commit"
Synchronizing configuration...OK.
Pica8 PicOS Version 2.5
Welcome to PicOS L2/L3 on XorPlus
admin@XorPlus>
Execute command: configure.
Entering configuration mode.
There are no other users in configuration mode.
admin@XorPlus#
Execute command: set vlans vlan-id 100.
admin@XorPlus#
Execute command: commit
.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus$

```

Show the Configuration via Linux Shell

```

admin@XorPlus$cli -c "show vlans"
Synchronizing configuration...OK.
Pica8 PicOS Version 2.5
Welcome to PicOS L2/L3 on XorPlus
admin@XorPlus>
Execute command: show vlans
.
VlanID   Tag           Interfaces
-----
1        untagged      ge-1/1/1, ge-1/1/2, ge-1/1/3, ge-1/1/4, ge-1/1/5
          ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9, ge-1/1/10
          ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14, ge-1/1/15
          ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19, ge-1/1/20
          ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24, ge-1/1/25
          ge-1/1/26, ge-1/1/27, ge-1/1/28, ge-1/1/29, ge-1/1/30
          ge-1/1/31, ge-1/1/32, ge-1/1/33, ge-1/1/34, ge-1/1/35
          ge-1/1/36, ge-1/1/37, ge-1/1/38, ge-1/1/39, ge-1/1/40
          ge-1/1/41, ge-1/1/42, ge-1/1/43, ge-1/1/44, ge-1/1/45
          ge-1/1/46, ge-1/1/47, ge-1/1/48, te-1/1/49, te-1/1/50
          te-1/1/51, te-1/1/52
          tagged
100      untagged
          tagged
admin@XorPlus>
admin@XorPlus$

```

Bash Linux Shell

User can execute Linux commands in the PicOS CLI to display the system process, create a directory, or execute commands added by third party software.

```

admin@XorPlus# run bash "ps"
  PID TTY          TIME CMD
 5289 ttyS0      00:00:00 bash
 5301 ttyS0      00:00:03 pica_sh
 7725 ttyS0      00:00:00 ps
admin@XorPlus# run bash "pwd"
/tmp/home/admin
admin@XorPlus>

```

If the command requires multiple parameters, quotation marks are required. Here is an example from the configuration mode used to check the system configuration file.

```
TelAviv# run bash "cat /pica/config/pica.conf"
/*XORP Configuration File, v1.0*/
/* Copyright (C) 2009-2013 Pica8, Inc.*/
/*Last commit : Fri May 9 12:43:39 2014 by admin*/
/*PicOS Version : 2.3*/
/*Version Checksum: 6b1435290092ce1b89fb98c06e20e66c*/
[...]
```

PicOS Upgrade and Configuration Change

As part of PicOS improvements, the CLI configuration structure may change between releases. Some commands or knobs may be added, removed, or modified. At the same time, a switch should be able to be upgraded without impact to a network, while still keeping the configuration intact.

A process was defined to achieve both of those goals (improved CLI with smooth upgrade).

When a command is modified in a release, the old command is marked with a "deprecated" flag.

A "deprecated" command can still be used on the next PicOS version, and, at a minimum, on all the PicOS versions published during the life time of this version (typically 9 months for a standard release).

Deprecated commands are hidden from the CLI, so a new user cannot access it without specific knowledge of the command (old commands cannot be auto-completed by tab key or shown by "?"). The old command must be fully entered manually to be used.

A warning message will be shown when a **commit** is done using a deprecated command.

Here is an example:

On PicOS 2.4, the command of "set interface management-ethernet" was deprecated and replaced by "set system management-ethernet."

This means that both of these commands work on PicOS 2.4, and the "interface management-ethernet" will be removed in a future release (9 months after PicOS 2.4 is published).

When using "set interface management-ethernet" in CLI, user will get the following information:

Configure node "interface management-ethernet" has been deprecated in version 2.4, please use "system management-ethernet" instead.



Note: The prompted information will not disappear until you remove the deprecated command.

When upgrading an image from an old version to a new one with configuration save, there will be some configuration nodes that are marked as deprecated in the new version, and user will get some notice information when committing in CLI.

For example, the following example reflects what may have happened during an upgrade from 2.3 to 2.4 with configuration saved, after upgrading and removing VLAN:

```

admin@XorPlus# delete vlans vlan-id 111
Deleting:
111 {
}
OK
admin@XorPlus# commit
Commit OK.
Configure node "interface management-ethernet" has been deprecated in version 2.4, please
use "system management-ethernet" instead.
Configure node "system syslog host" has been deprecated in version 2.4, please use "system
syslog server-ip" instead.
Configure node "system syslog port-number" has been deprecated in version 2.4, please use
"system syslog server-ip" instead.
Configure node "system syslog port-protocol" has been deprecated in version 2.4, please use
"system syslog server-ip" instead.
Save done.
admin@XorPlus#

```



When upgrading image from an old version to a new one, with a configuration that contains deprecated nodes, the upgrade will fail. Remove the deprecated nodes on the configuration tree. Then, install upgrade again.

Set CLI

We support setting some values for CLI.

```

admin@XorPlus> set cli ?
Possible completions:
idle-timeout Set maximum idle time before login session ends
screen-length Set terminal screen length
terminal Set terminal type
admin@XorPlus> set cli screen-length ?
Possible completions:
<length> Number of lines of text that the terminal screen displays (0..10,000).
admin@XorPlus> set cli ?
Possible completions:
idle-timeout Set maximum idle time before login session ends
screen-length Set terminal screen length
terminal Set terminal type
admin@XorPlus> set cli idle-timeout 0
admin@XorPlus> set cli screen-length ?
Possible completions:
<length> Number of lines of text that the terminal screen displays (0..10,000).
admin@XorPlus> set cli screen-length 0
admin@XorPlus>
admin@XorPlus> set cli terminal ?
Possible completions:
ansi ANSI-compatible terminal
linux Linux-compatible terminal
vt100 VT100-compatible terminal
xterm Xterm window terminal
admin@XorPlus> set cli terminal

```

System Administration

- Configuring DHCP and Static IP Address
- Configuring DHCP Relay
- Configuring DHCP Relay after PicOS 2.9.2(including 2.9.2)
- Configuring DHCP Option82
- Configuring DHCP Snooping
- Configuring DHCP Snooping after PicOS 2.9.2(including 2.9.2)
- Configuring DHCP Relay and DHCP Snooping together after PicOS 2.9.2(including 2.9.2)
- Configuring a User Account
- Configuring Authentication, Authorization, and Accounting
- Configuring SSH and Telnet Parameters
- Configuring the Log-in ACL
- Configuring NTP and the Time Zone Parameter
- Configuring PTP
- Configuring the linux-config-unreliable
- Configuring IPFIX
- Configuring sFlow
- Configuring SNMPv2
- Configuring the Syslog Log Level
- Configuring the Syslog Disk and Syslog Host
- Displaying System Information
- Technical Support
- Flushing ARP and the Neighbor Table
- Rebooting the System
- Displaying the Debugging Message
- System Management Command List
- IPv6 Management Support
- PoE Configuration
- NETCONF Configuration
- Configuring SNMPv3

Configuring DHCP and Static IP Address

This document describes the process to configure the management IP address on a PicOS switch. The switch can obtain its IP address through DHCP (Dynamic Host Configuration Protocol), or user can assign a static IP address.

Enabling DHCP

By default, DHCP is enabled on the management interface **eth0**. User can manually enable DHCP with the following command in PicOS L2/L3 configuration mode.

```
admin@XorPlus# set system management-ethernet eth0 ip-address IPv4 dhcp
admin@XorPlus# com
Commit OK.
Save done.
admin@XorPlus#
```

Configuring a Static IPv4 Address

User can configure the switch management interface **eth0** with a static IPv4 address and default gateway, as shown below.

```
admin@XorPlus# set system management-ethernet eth0 ip-address IPv4 10.10.50.139/24
admin@XorPlus# set system management-ethernet eth0 ip-gateway IPv4 10.10.50.1
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring a Static IPv6 Address

User can configure the switch management interface **eth0** with a static IPv6 address and default gateway, as shown below.

```
admin@XorPlus# set system management-ethernet eth0 ip-address IPv6 2001::123/64
admin@XorPlus# set system management-ethernet eth0 ip-gateway IPv6 2001::1
admin@XorPlus# commit
Commit OK.
Save done.
```

Enabling DHCP Relay in a VLAN Interface

When you enable the DHCP relay in a VLAN interface, the switch will relay the received DHCP request to the specified DHCP server, via routing. Normally, the port that connects to a trusted DHCP server should be a trusted port. You should configure the port using the *trust true* option.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols dhcp relay vlan-interface vlan-2 disable false
admin@XorPlus# set protocols dhcp relay vlan-interface vlan-2 dhcp-server-address1
192.168.2.100
admin@XorPlus# set protocols dhcp snooping port ge-1/1/2 trust true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configuring DHCP Relay after PicOS 2.9.2(including 2.9.2)

As broadcast is used to send the request message during the dynamic acquisition of IP address, DHCP is only applicable to the case where the DHCP client and the server are on the same subnet. For dynamic host configuration, you need to set up a DHCP server on all network segments, which is obviously very uneconomical.

The introduction of DHCP relay function solves this problem: the client can communicate with the DHCP server of other network segments through DHCP relay, and finally obtains the IP address. In this way, DHCP clients on multiple networks can use the same DHCP server, saving cost and facilitating centralized management.

Procedure

Step 1. Configure VLAN.

a). Configure VLAN ID.

set vlans <vlan-id>

b). Add an interface to the VLAN.

set interface gigabit-ethernet <port> **family ethernet-switching native-vlan-id** <vlan-id>

c). Configure the access/trunk mode.

set interface gigabit-ethernet <port> **family ethernet-switching port-mode** <port-mode>

d). Configure the IP address for the VLAN interface to implement Layer 3 connectivity .

```
set vlan-interface interface <interface-name> vif <vif-name> address <address> prefix-length <number>
```

Step 2. Associate a Layer 3 interface with a VLAN.

```
set vlans vlan-id <vlan-id> I3-interface <interface-name>
```

Step 3. Enable the DHCP relay function on the VLAN interface.

```
set protocols dhcp relay vlan-interface <interface-name> disable {true | false}
```

Step 4. Configure the IP address of the DHCP server.

```
set protocols dhcp relay vlan-interface <vlan-interface-name> { dhcp-server-address1 address | dhcp-server-address2 address | dhcp-server-address3 address | dhcp-server-address4 address }
```

Step 5. Configure the interface connected to the DHCP server as DHCP snooping trusted interface.

```
set protocols dhcp snooping port <interface-name> trust {true | false}
```

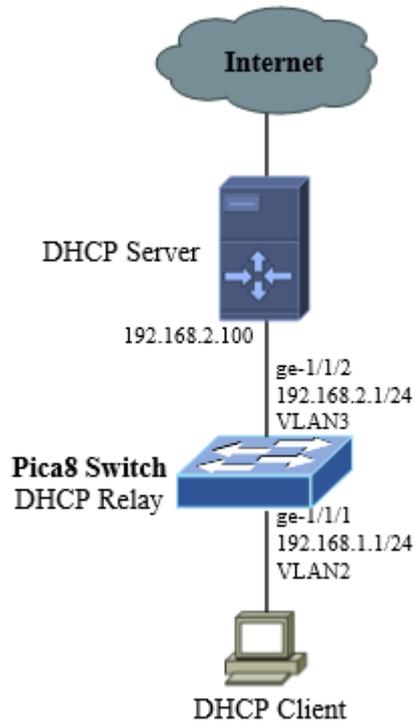
NOTE

- If the interface connected to the DHCP server is not configured as a DHCP snooping trusted interface, the data packets sent from the DHCP server to the host will be dropped, the DHCP function will not work.
- On the same switch, the DHCP client and DHCP relay functions are mutually exclusive. That is to say, if the command **set vlan-interface interface** <interface-name> **dhcp true** is configured on a switch, setting a VLAN interface as a DHCP client, then enabling the DHCP relay function on the same device will have the net effect of the DHCP replay function being enabled whereas the client function will not be enabled.

Configuration Example

Networking Requirements

- The IP address of the ge-1/1/1 interface of the Pica8 Switch is 192.168.1.1/24 and belongs to VLAN 2. The IP address of the ge-1/1/2 interface is 192.168.2.1/24 and belongs to VLAN 3.
- Enable the DHCP relay function on VLAN 2.
- The IP address of DHCP server is 192.168.2.100.



Procedure

Step 1. Configure VLAN.

```
admin@XorPlus# set vlans vlan-id 2
```

```
admin@XorPlus# set vlans vlan-id 3
```

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 2
```

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 3
```

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode trunk
```

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode trunk
```

```
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1 prefix-length 24
```

```
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1 prefix-length 24
```

```
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
```

```
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
```

Step 2. Enable the DHCP relay function on VLAN 2.

```
admin@XorPlus# set protocols dhcp relay vlan-interface vlan-2 disable false
```

Step 3. Configure the IP address of DHCP server as 192.168.2.100.

```
admin@XorPlus# set protocols dhcp relay vlan-interface vlan-2 dhcp-server-address1 192.168.2.100
```

Step 4. Configure the interface connected to the DHCP server as DHCP snooping trusted interface.

```
admin@XorPlus#set protocols dhcp snooping port ge-1/1/2 trust true
```

Step 5. Commit the configuration.

```
admin@XorPlus# commit
```

Step 6. Verify the configuration.

After the configuration is complete, run the **show protocols dhcp relay** command to view the configuration of DHCP relay.

```
admin@Xorplus# show protocols dhcp relay
```

```
relay {  
vlan-interface vlan2 {  
dhcp-server-address1: 192.168.2.100  
}  
}
```

Configuring DHCP Option82

Option82 is a relay agent used to specify the DHCP client location information. The DHCP Option82 is disabled by default. To enable Option82, use the *disable false* option. Then, use the *circuit-id* command to set the DHCP port information.

Enable DHCP Option82

```
admin@XorPlus# set protocols dhcp option82 disable false  
admin@XorPlus# commit  
Merging the configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

Modify the Circuit-id of Option82

```
admin@XorPlus# set protocols dhcp relay port ge-1/1/3 circuit-id v100  
admin@XorPlus# commit  
Merging the configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring DHCP Snooping

DHCP snooping creates a mapping table, which includes the IP address, the MAC address, and the port number. DHCP snooping is disabled by default. The steps below explain how to enable DHCP snooping and configure the DHCP snooping binding file, trust port (by default the port is untrusted), and timeout functions.

Enable DHCP Snooping

```
admin@XorPlus# set protocols dhcp snooping disable false
admin@XorPlus# commit
Commit OK.
Save done.
XorPlus#
```

Configure DHCP Snooping Binding File and Timeout

Sync the DHCP snooping table to disk.

```
admin@XorPlus# set protocols dhcp snooping binding file /tmp/run/dhcp_bind
admin@XorPlus# set protocols dhcp snooping binding timeout 8
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure DHCP Snooping Trust Port

DHCP reply is usually trusted. The port connected to the DHCP server should enable this.

```
admin@XorPlus# set protocols dhcp snooping port ge-1/1/2 trust true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Display the DHCP Snooping Table of Host Information

```
admin@XorPlus# run show dhcp snooping
Total count: 1
MAC Address IP Address Port VLAN ID VLAN Interface
-----
00:1d:09:fa:a1:b4 192.168.1.10 ge-1/1/1 2 vlan2
```

Configuring DHCP Snooping after PicOS 2.9.2(including 2.9.2)

DHCP snooping creates a mapping table, which includes the IP address, the MAC address, and the port number. DHCP snooping is disabled by default. The steps below explain how to enable DHCP snooping and configure the DHCP snooping binding file, trust port (by default the port is untrusted), and timeout functions.

Procedure

Step 1. Enable DHCP snooping on a VLAN.

```
set protocols dhcp snooping vlan {<vlan-id> | all}
```

Step 2. Configure the interface connected to the DHCP server as DHCP snooping trusted interface.

```
set protocols dhcp snooping port <interface-name> trust {true | false}
```

NOTE

The DHCP snooping trusted interface can not be a LAG interface.

Step 3. (Optional) Configure the DHCP snooping binding file and the value of timeout.

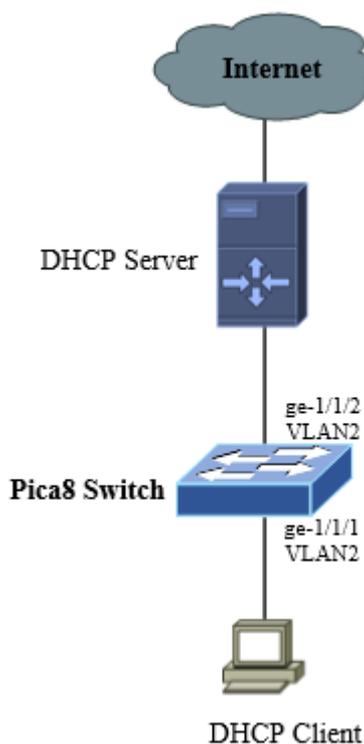
```
set protocols dhcp snooping binding file <file>
```

```
set protocols dhcp snooping binding timeout <time>
```

Configuration example

Networking Requirements

- Configure the ge-1/1/1 interface and ge-1/1/2 interface to VLAN 2.
- Enable DHCP snooping on VLAN 2 and configure the interface connected to the DHCP server as the DHCP snooping trusted interface.



Procedure

Step 1. Configure VLAN.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 2
```

Step 2. Enable DHCP snooping on VLAN 2.

```
admin@XorPlus#set protocols dhcp snooping vlan 2
```

Step 3. Configure the interface connected to the DHCP server as DHCP snooping trusted interface.

```
admin@XorPlus# set protocols dhcp snooping port ge-1/1/2 trust true
```

Step 4. (Optional) Configure /tmp/run/dhcp_bind as the DHCP snooping binding file and the value of timeout is 8.

```
admin@XorPlus# set protocols dhcp snooping binding file /tmp/run/dhcp_bind
admin@XorPlus# set protocols dhcp snooping binding timeout 8
```

Step 5. Commit the configuration.

```
admin@XorPlus# commit
```

Step 6. Verify the configuration.

- • • After the configuration is complete, run the **show protocols dhcp snooping** command to view the configuration of DHCP snooping.

```
admin@Xorplus# show protocols dhcp snooping
snooping {
disable: false
}
```

- • • After the configuration is complete, run the **run show dhcp snooping** command to view the DHCP snooping binding table.

```
admin@XorPlus# run show dhcp snooping
Total count: 1
MAC Address IP Address Port VLAN ID VLAN Interface
-----
```

```
00:1d:09:fa:a1:b4 192.168.1.10 ge-1/1/1 2
```

The VLAN Interface has value when DHCP relay is configured to the VLAN interface, otherwise the value is null.

Configuring DHCP Relay and DHCP Snooping together after PicOS 2.9.2(including 2.9.2)

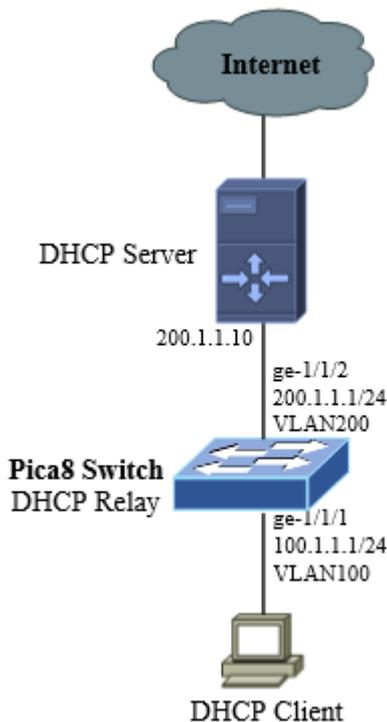
By configuring DHCP relay and DHCP snooping together for ARP snooping defense, about DHCP relay configuration please refer to Configuring DHCP Relay after PicOS 2.9.2(including 2.9.2), about DHCP snooping configuration please refer to Configuring DHCP Snooping after PicOS 2.9.2(including 2.9.2).

If we configure the DHCP relay and DHCP snooping together, the mapping table of DHCP snooping will be synchronized to ARP inspection table to validate ARP packets in a network, please refer to Dynamic ARP Inspection for detail of ARP inspection table .

Configuration example

Networking Requirements

- The IP address of the ge-1/1/1 interface of Pica8 Switch is 100.1.1.1/24 and belongs to VLAN 100, and enable DHCP relay function on VLAN 100.
- The IP address of the ge-1/1/2 interface of Pica8 Switch is 200.1.1.1/24 and belongs to VLAN 200, and configure the interface as the DHCP snooping trusted interface.
- Enable DHCP snooping on VLAN 100.
- The IP address of DHCP server is 200.1.1.10.



Procedure

Step 1. Configure VLAN.

```
admin@XorPlus#set vlans vlan-id 100
admin@XorPlus#set vlans vlan-id 200
admin@XorPlus#set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 100
admin@XorPlus#set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 200
admin@XorPlus#set vlan-interface interface vlan-100 vif vlan-100 address 100.1.1.1 prefix-length 24
admin@XorPlus#set vlan-interface interface vlan-200 vif vlan-200 address 200.1.1.1 prefix-length 24
admin@XorPlus#set vlans vlan-id 100 l3-interface vlan-100
admin@XorPlus#set vlans vlan-id 200 l3-interface vlan-200
```

Step 2. Enable DHCP relay on VLAN 100.

```
admin@XorPlus#set protocols dhcp relay vlan-interface vlan-100 disable false
```

Step 3. Configure the IP address of DHCP server as 200.1.1.10.

```
admin@XorPlus#set protocols dhcp relay vlan-interface vlan-100 dhcp-server-address1 200.1.1.10
```

Step 4. Enable DHCP snooping on VLAN 100.

```
admin@XorPlus#set protocols dhcp snooping vlan 100
```

NOTE:

The VLAN that enabling DHCP snooping needs to be configured as the VLAN to which the interface connected to the host.

Step 5. Configure the interface ge-1/1/2 as DHCP snooping trusted interface.

```
admin@XorPlus#set protocols dhcp snooping port ge-1/1/2 trust true
```

Step 6. Commit the configuration.

```
admin@XorPlus# commit
```

Step 7. Verify the configuration.

- • • After the configuration is complete, run the **show protocols dhcp** command to view the configuration.

```
admin@Xorplus# show protocols dhcp
```

```
relay {
```

```
vlan-interface vlan-100 {
```

```
dhcp-server-address1: 200.1.1.10
```

```
}
```

```
}
```

```
snooping {
```

```
disable: false
```

```
}
```

- • • After the configuration is complete, run the **run show dhcp snooping** command to view the information of DHCP snooping binding table.

```
admin@XorPlus# run show dhcp snooping
```

```
Total count: 1
```

```
MAC Address IP Address Port VLAN ID VLAN Interface
```

```
-----
```

```
00:1d:09:fa:a1:b4 100.1.1.11 ge-1/1/1 100 vlan-100
```

Configuring a User Account

There are two types of user accounts: **super-user** and **read-only**. The newly created user account, by default, is read-only.

Creating a User Class and Password

```
admin@XorPlus# set system login user ychen authentication plain-text-password pica8
admin@XorPlus#set system login user ychen class super-user
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

Configuring a T elnet Announcement

```
admin@XorPlus# set system login announcement "welcome the switch-1101"
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Authentication, Authorization, and Accounting

PicOS supports Authentication/Authorization/Accounting (AAA). Once authenticated by the AAA server (referred to as "admin" in our guide), user can configure the switch. PicOS supports TACACS+ and RADIUS protocols. RADIUS supports only two levels: **read-only** and **super-user**.

Configure the local switch and server as shown below:

Configuring AAA in the Switch

Configure the tacacs enable:

```
admin@XorPlus# set system aaa tacacs-plus disable false
admin@XorPlus# set system aaa tacacs-plus key pica8
admin@XorPlus# set system aaa tacacs-plus server-ip 10.10.53.53
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set system aaa tacacs-plus authorization true
admin@XorPlus# set system aaa tacacs-plus accounting true
admin@XorPlus# commit
```

Configure the radius enable:

```
admin@XorPlus# set system aaa radius authorization disable false
admin@XorPlus# set system aaa radius authorization server-ip 10.10.50.41 shared-key
testing123
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set system aaa radius accounting disable false
admin@XorPlus# set system aaa radius accounting server-ip 10.10.50.41 shared-key testing123
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
XorPlus#
```

Displaying AAA information

```
admin@XorPlus# show system aaa tacacs-plus
Waiting for building configuration.
authorization: true
accounting: true
server-ip 10.10.53.53
key: "pica8"
XorPlus# show system aaa radius
Building the configuration.
authorization {
disable: false
server-ip 10.10.50.41 {
shared-key: "testing123"
}
}
accounting {
disable: false
server-ip 10.10.50.41 {
shared-key: "testing123"
}
}
admin@XorPlus#
```

Configuring the AAA Server

Configure the AAA server configuration file as follows:

Tacacs server configuration:

key = pica8

1. Accounting File

```
accounting file = /var/tmp/acctfile
default authentication = file /etc/passwd
user = admin {
  member = admins
}
group = admins {
  global = cleartext "password"
  service = exec {
    default attribute = permit
  }
}
user = operator {
  global = cleartext "operator"
  service = exec {
    default attribute = permit
  }
}
user = ychen {
  global = cleartext "ychen"
  member = admins
  service = exec {
    default attribute = permit
  }
}
Add "/"
```

usr/share/freeradius/dictionary.pica8" to radius server before the configuration.

Radius server configuration:

```
operator Cleartext-Password := "testing"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Framed-IP-Address = 172.16.3.33,
Framed-IP-Netmask = 255.255.255.0,
Framed-Routing = Broadcast-Listen,
Framed-Filter-Id = "std.ppp",
Framed-MTU = 1500,
Framed-Compression = Van-Jacobsen-TCP-IP,
Class = "read-only"
ychen Cleartext-Password := "testing"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Framed-IP-Address = 172.16.3.33,
Framed-IP-Netmask = 255.255.255.0,
Framed-Routing = Broadcast-Listen,
Framed-Filter-Id = "std.ppp",
Framed-MTU = 1500,
Framed-Compression = Van-Jacobsen-TCP-IP,
Class = "super-user"
```

Following the configuration above, the admin or operator can access the switch via telnet or SSH. Any valid CLI commands executed by the admin or operator will be recorded to the specified accounting file. In our example above, the accounting file is */var/tmp/acctfile*.

Configuring the Local Log-in

```
admin@XorPlus# set system aaa local disable true
admin@XorPlus# commit
Commit OK.
Save done.
```

In the configuration above, user cannot log in to the switch with a local account.

Configuring SSH and Telnet Parameters

Configuring the SSH Connection Limit

```
admin@XorPlus# set system services ssh protocol-version v2
admin@XorPlus# set system services ssh connection-limit 5
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Disabling Telnet Service

```
admin@XorPlus# set system services telnet disable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Enabling and Disabling Inband Service

By default, SSH and telnet with inband interfaces are disabled. You can enable inband services by entering the command below.

```
admin@XorPlus# set system inband enable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the Log-in ACL

Configure the ACL to control whether remote hosts within specified sub-networks are allowed to log in to the system. In our example, remote hosts from both sub-networks that we configured may log in.

```
admin@XorPlus# set system login-acl network 192.168.1.0/24
admin@XorPlus# set system login-acl network 192.168.100.100/32
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring NTP and the Time Zone Parameter

Configuring the NTP Server IP Address

The L2/L3 switch synchronizes with the NTP server only when the configuration commands are committed using the **commit** command. User can change the NTP server's IP address, as shown below.

```
admin@XorPlus# set system ntp-server-ip 192.168.10.100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the Time Zone

Configure the time zone as follows (The time zone configured below is Pacific/Kosrae).

```
admin@XorPlus# set system timezone Pacific/Kosrae
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the System Clock

```
admin@XorPlus> set date 2012.01.01-23:59
Sun Jan 1 23:59:00 UTC 2012
admin@XorPlus>
```

The clock will be set in the hardware.

Configuring PTP

- Principle
 - Introduction to PTP
 - Basic Concepts
 - Delay Request-Response Mechanism
- Product Support
- PTP Configuration
 - Procedure
- Configuration Example

Principle

 The PicOS supports PTP function since 2.9.2.

Introduction to PTP

On a modern communication network, most telecommunications services require that the frequency offset or time difference between devices is within an acceptable range. To meet this requirement, network clock synchronization must be implemented.

Network clock synchronization includes phase synchronization and frequency synchronization.

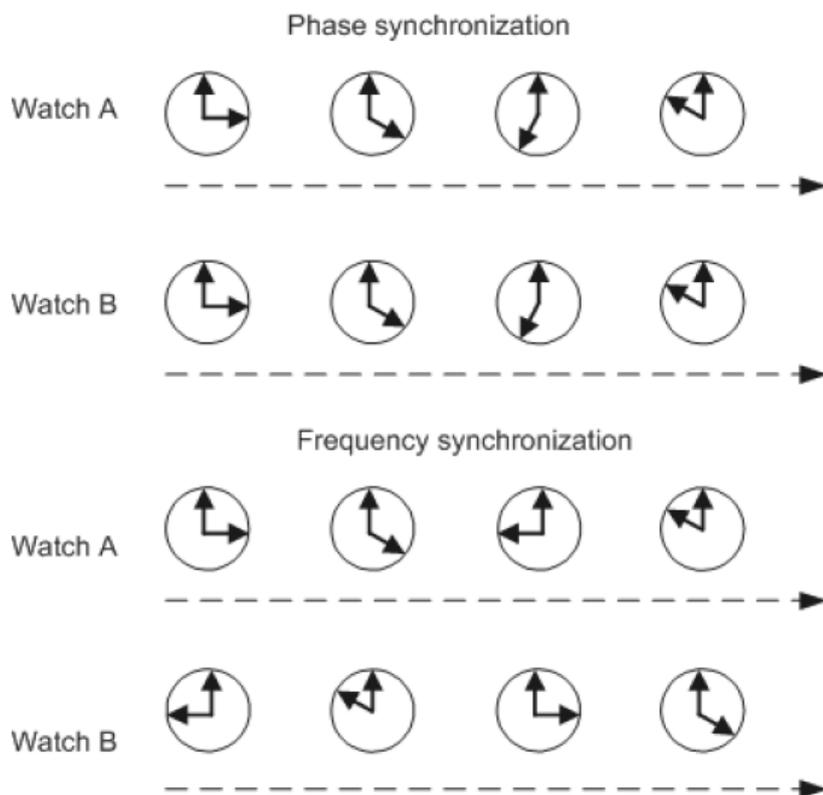
Phase synchronization

Phase synchronization, also called time synchronization, refers to the consistency of both frequencies and phases between signals. That is to say, the phase difference between signals is always 0.

Frequency synchronization

Frequency synchronization refers to constant phase difference between signals. It ensures that signals are sent or received at the same rate at a moment so that all devices on the communication network operate at the same rate.

Figure 1 Difference between phase synchronization and frequency synchronization



PTP (Precision Time Protocol) is a time synchronization protocol, which is not only used for high-precision time synchronization between devices, but can also be used to synchronize the frequency between devices. Compared to the existing time synchronization mechanism, PTP has the following advantages:

- Compared to NTP (Network Time Protocol), PTP can meet the more accurate time synchronization requirements, NTP generally can only achieve sub-second time synchronization accuracy, and PTP can reach sub-microsecond.
- Compared to GPS (Global Positioning System), PTP has lower construction and maintenance costs, and because it can get rid of the dependence on GPS, it has special meaning in national security.

Basic Concepts

PTP domain

A PTP domain is a logical grouping of clocks that synchronize to each other using the protocol IEEE 1588v2, but that are not necessarily synchronized to clocks in another domain. Each PTP domain is an independent PTP clock synchronization system and has only one clock source.

Clock node

Clock nodes are nodes in a PTP domain. PTP defines the following types of clock nodes:

- Ordinary clock (OC) device: provides only one physical port to participate in time synchronization in a PTP domain. An OC device uses this port to synchronize time with an upstream device or send time to a downstream device.
- Boundary clock (BC) device: provides two or more physical ports to participate in time synchronization in a PTP domain. One port synchronizes time with an upstream device, and the others send the time to downstream devices. A clock node is also a BC device if it functions as the clock source and sends time to downstream devices through multiple PTP ports.
- Transparent clock (TC) device: forwards PTP messages between its PTP ports and measures the link delay of the messages. Different from an OC device and a BC device, a TC device does not synchronize time with other devices through ports.
 - E2ETC (End-to-End Transparent Clock): forwards packets of non-P2P (non-Peer-to-Peer) types directly on the network and participates in the calculation of the entire link Time.

E2ETC calculate the residence time of the device, which is the time it takes the event message from ingress port to egress port, then add the calculated residence time to the correctionfield of the event message.

- P2PTC (Peer-to-Peer Transparent Clock): forwards only Sync messages, Follow_Up messages and Announce messages, and terminates other PTP packets, and participates in the calculation the delay of each link on the entire link.

PTP port

A PTP port is a port running PTP. PTP ports are classified into the following types based on roles:

- Master port: The port is the source of time on the path served by the port, located on a BC or OC device.
- Slave port: The port synchronizes to the device on the path with the port that is in the MASTER state, located on a BC or OC device.
- Passive port: The port is not the master on the path nor does it synchronize to a master. It is an idle port on a BC device and does not receive or send synchronization clock signals.

Master-slave hierarchy

Nodes in a PTP domain establish the master-slave hierarchy for clock synchronization. Master nodes send synchronization clock signals, while slave nodes receive synchronization clock signals. A device may receive synchronization clock signals from an upstream node and then send the synchronization clock signals to a downstream device.

If two clock nodes synchronize time with each other:

- The node that sends synchronization clock signals is the master node, and the node that receives synchronization clock signals is the slave node.

- The clock on the master node is the master clock, and the clock on the slave node is the slave clock.
- The port that sends synchronization clock signals is the master port, and the port that receives synchronization clock signals is the slave port.

Grandmaster clock

All clock nodes in a PTP domain are organized into the master-slave hierarchy. The grandmaster clock (GMC) is at the top of the hierarchy and is the reference clock in the PTP domain. Clock nodes exchange PTP messages to synchronize the time of the GMC to the entire PTP domain. Therefore, the GMC is also called the clock source. The GMC can be statically configured or dynamically elected through the best master clock (BMC) algorithm.

PTP message

Nodes exchange PTP messages to establish the master-slave hierarchy and implement time and frequency synchronization. PTP messages are classified into event messages and general messages depending on timestamps:

- Event message: is tagged with a timestamp when reaching or leaving a port. PTP devices calculate the link delay based on the timestamps carried in event messages. Event messages include Sync, Delay_Req, Pdelay_Req, and Pdelay_Resp messages.
- General message: is used to establish master-slave hierarchy, and to request and send time information. General messages are not tagged with timestamps. General messages include Announce, Follow_Up, Delay_Resp, Pdelay_Resp_Follow_Up, Management, and Signaling messages.

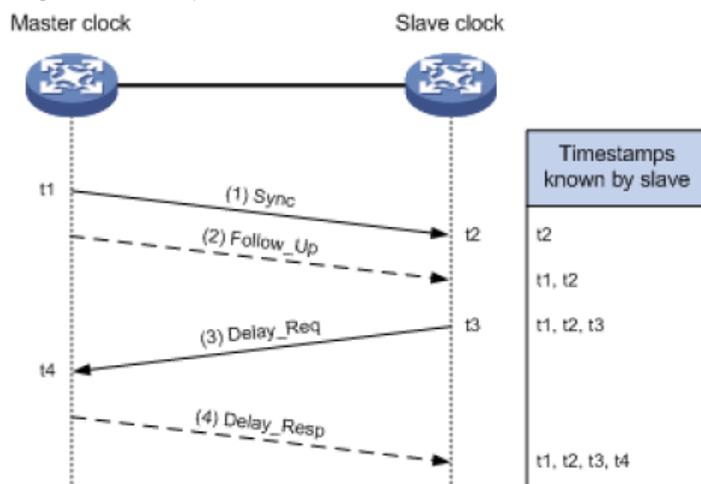
NOTE

Currently, devices supports only E2ETC function, and uses Sync, Delay_Req and Delay_Resp PTP messages.

Delay Request-Response Mechanism

Figure 2 shows the process of calculating the average link delay and time offset between the master and slave devices using the E2E mechanism.

Figure 2 Delay request-response mechanism



1. The master sends a Sync message to the slave and notes the time t1 at which it was sent.
2. The slave receives the Sync message and notes the time of reception t2.
3. The master conveys to the slave the timestamp t1 by:
 - a) Embedding the timestamp t1 in the Sync message. This requires some sort of hardware processing for highest accuracy and precision.
 - b) Embedding the timestamp t1 in a Follow_Up message.
4. The slave sends a Delay_Req message to the master and notes the time t3 at which it was sent.
5. The master receives the Delay_Req message and notes the time of reception t4.

The master conveys to the slave the timestamp t4 by embedding it in a Delay_Resp message. By exchanging messages with the master device, the slave device obtains t1, t2, t3, and t4, and calculates the average link delay and time offset between the master and slave devices. Then the slave device can adjust the local time according to the calculated time offset to synchronize with the master device. The formulas for calculating the link delay and time offset are as follows:

$$\text{Average link delay} = [(t4 - t1) - (t3 - t2)]/2$$

$$\text{Offset} = [(t2 - t1) + (t3 - t4)]/2$$

Product Support

The table below shows the model and the corresponding switch ASIC that support the PTP function, others do not support.

Switch ASIC		Model
Helix4	BCM56342	as4610-30t
		as4610_30p
	BCM56340	as4610_54t
		as4610_54p
Trident2	BCM56850	pronto5101
		pronto5401
		as6701-32x
		as6712-32x
	BCM56854	pronto5101
		as5712-54x
		s4048
		arctica4806xp
Trident2Plus	BCM56864	as5812-54t
		as5812-54x

Tomahawk	BCM56960	dc57032q28
		z9100
		as7312
		as7712-32x

PTP Configuration

Configure the device interface as the E2ETC node type, enable the PTP function on the interface, and implement the device to forward the PTP packets directly.

Procedure

Step 1. Configure the device interface as the E2ETC node type, enable the PTP function on the interface.

```
set interface gigabit-ethernet <port> ptp mode {e2etransparent | none}
```

By default, PTP is disabled on an interface.

NOTE

- The PTP function can only be configured on the physical interface.
- On the TC device, PTP must be enabled on all the interfaces that receive and send PTP messages. Otherwise, the PTP function will work improperly.

Configuration Example

Procedure

Step 1. Configure the port te-1/1/25 that receiving PTP messages and the port te-1/1/26 that sends PTP packets as E2ETC node.

```
admin@Xorplus#set interface gigabit-ethernet te-1/1/25 ptp mode e2etransparent
```

```
admin@Xorplus#set interface gigabit-ethernet te-1/1/26 ptp mode e2etransparent
```

Step 2. Verify the configuration.

After the configuration is complete, run the **run show interface detail** command to view the configuration of PTP.

```

admin@Xorplus# run show interface gigabit-ethernet te-1/1/25 detail
Physical interface: te-1/1/25, Enabled, error-discard False, Physical link is Down
Interface index: 25, Mac Learning Enabled
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit ingress:0, egress:0
Precision Time Protocol mode:e2etransparent
Current address: 48:6e:73:01:00:61, Hardware address: 48:6e:73:01:00:61
Traffic statistics:
5 sec input rate 0 bits/sec, 0 packets/sec
5 sec output rate 0 bits/sec, 0 packets/sec
-More-

```

Configuring the linux-config-unreliable

PicOS is a Linux distribution with Pica8 tools for Routing, Switching, and OpenFlow. As such, PicOS can be configured directly from the Linux shell.

But sometimes, operators would like to be sure the CLI configuration completely reflects the state of the system. It can be useful if the same parameter can be configured in the CLI, or in Linux at the same time. A good example of such parameters would be the IP address of the system or its default gateway.

In the PicOS CLI configuration, all of those parameters have been gathered under the system hierarchy. It is then the choice of the operator to choose either the CLI or Linux shell to configure the system parameters.

A new knob has been added for control, in case the CLI configuration should override the Linux configuration.

By default, the CLI configuration does NOT override the Linux system configuration.



When the linux-config-unreliable knob is modified, please commit it first, then set other system settings. Do not commit the linux-config-unreliable and system delta at the same time.

You can choose bash control or xorp control. As shown below:

Xorp Control

```

admin@XorPlus# set system linux-config-unreliable true
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# show system
  linux-config-unreliable: true
  services {
    telnet {
      disable: false
    }
  }
  log-level: "trace"

```

Bash Control

```
admin@XorPlus# set system linux-config-unreliable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# show system
  linux-config-unreliable: false
  log-level: "trace"
```

If user chooses bash control, the system settings should be set in bash, not from xorp. Otherwise, the system command should be set in xorp.

```
admin@XorPlus# show system
  linux-config-unreliable: false
  log-level: "trace"
admin@XorPlus# set system hostname pica8
admin@XorPlus# commit
The system is managed by linux
Commit failed.
admin@XorPlus#
```



The following system commands should always be set in PicOS, even when system is under bash control:

```
system aaa tacacs-plus
system log-level
```



When changing the system control straight from xorp control to bash control, the xorp configurations related to system will be removed from the configuration tree automatically. When changing from Bash control to xorp control, the configuration related to the system will be read and added into the xorp configuration tree automatically.

```
system log-facility
```

Configuring IPFIX Parameters

By default, IPFIX is disabled. You can enable IPFIX and configure its parameters, as shown below. Make sure the switch can connect to the IPFIX collector server correctly.

```
admin@XorPlus# set protocols ipfix collector 192.168.2.10 udp-port 9999
admin@XorPlus# set protocols ipfix interfaces ingress ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring sFlow

Globally Enabling sFlow

By default, sFlow is disabled. User can enable sFlow and configure the parameters, verify that the switch can connect to the sFlow collector server, and configure the sFlow **agent-id** and **source-address** at the same time that sFlow is enabled.

```
admin@XorPlus# set protocols sflow disable false
admin@XorPlus# set protocols sflow agent-id 10.10.50.248
admin@XorPlus# set protocols sflow source-address 10.10.50.248
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring sFlow Parameters

User can configure global parameters for sFlow, including agent-id, collector IP, polling-interval, sampling-rate, and source-address.

```
admin@XorPlus# set protocols sflow agent-id 10.10.50.248
admin@XorPlus# set protocols sflow collector 10.10.50.221 udp-port 6343
admin@XorPlus# set protocols sflow polling-interval 30
admin@XorPlus# set protocols sflow sampling-rate ingress 2000
admin@XorPlus# set protocols sflow sampling-rate egress 2000
admin@XorPlus# set protocols sflow header-len 128
admin@XorPlus# set protocols sflow source-address 10.10.50.248
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show sflow
sFlow : Enabled
Agent ID : 10.10.50.248
Source Address : 10.10.50.248
Sample rate ingress: 1:2000
Sample rate egress : 1:2000
Polling interval : 30 seconds
Header Length : 128
admin@XorPlus#
admin@XorPlus# run show sflow collector
Collector address UDP-port No of Samples
-----
10.10.50.221 6343 5336
admin@XorPlus#
```

Configuring sFlow on a Specific Interface

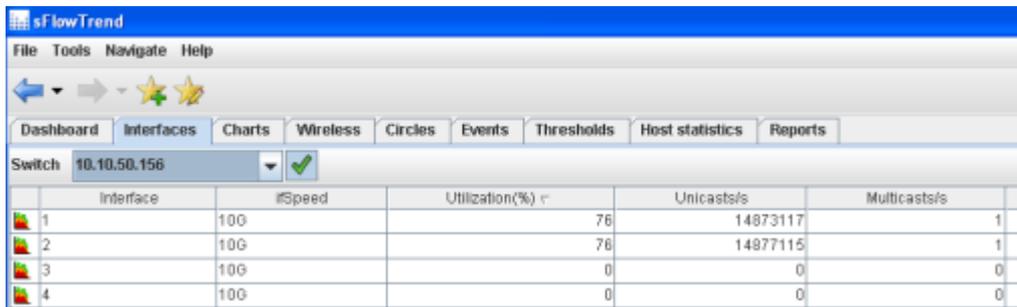
User can configure sFlow parameters on a specific interface but needs to enable sflow protocols on global first.

```

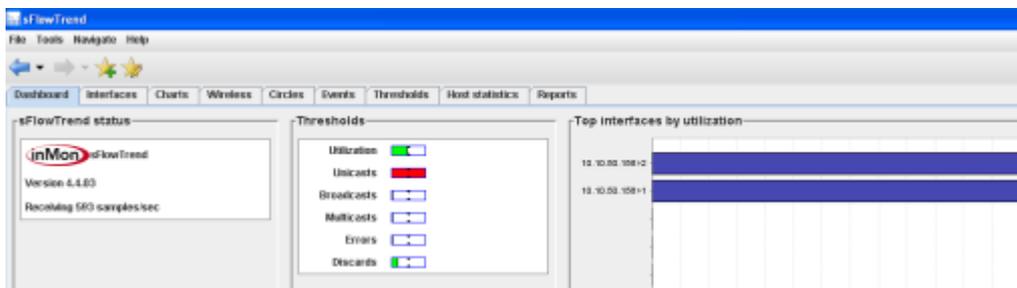
admin@XorPlus# set protocols sflow interface ge-1/1/1 ?
Possible completions:
  <[Enter]>           Execute this command
  disable            Disable sflow on all interfaces by default
  header-len        The Length of sampled packet in bytes, 64 by default
  polling-interval  How often the sflow agent polls the interface in seconds,
  30 by default
  sampling-rate     The rate at which packets must be sampled, 2000 by default
admin@XorPlus# set protocols sflow interface ge-1/1/1 disable false
admin@XorPlus# set protocols sflow interface ge-1/1/1 header-len 128
admin@XorPlus# set protocols sflow interface ge-1/1/1 polling-interval 10
admin@XorPlus# set protocols sflow interface ge-1/1/1 sampling-rate ingress 1000
admin@XorPlus# commit

```

In the current version, sFlow samples only the ingress traffic of each interface. User can monitor the traffic with sFlow Trend.



Interface	#Speed	Utilization(%) r	Unicasts/s	Multicasts/s
1	10G	76	14873117	1
2	10G	76	14877115	1
3	10G	0	0	0
4	10G	0	0	0



User needs to enable sFlow on global and on the port so that it can sample flow and counters. If sFlow is only enabled on global, it will only see the port count.

Configuring SNMPv2

Configuring SNMPv2 Parameters

By default, SNMP is disabled. You can enable SNMP and configure its parameters (e.g. community, contact, location).

```

admin@XorPlus# set protocols snmp community Pica8-data-center
admin@XorPlus# set protocols snmp community Pica8-data-center authorization read-only
admin@XorPlus# set protocols snmp contact support@pica8.com
admin@XorPlus# set protocols snmp location Beijing
admin@XorPlus# set protocols snmp trap-group targets 10.10.1.1 security-name
Pica8-data-center
admin@XorPlus# set protocols snmp trap-group version v2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```



In version 2.8.1, security-name has to be configured for trap-group targets whichever the version is .

Configuring an SNMP ACL

By default, all hosts can *snmpwalk* the information of the switch. Configure an SNMP ACL to control which hosts within the subnet can snmpwalk the switch.

```

admin@XorPlus# set system snmp-acl network 1.1.1.0/24
admin@XorPlus# set system snmp-acl network 2.2.2.0/24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring SNMPset

User can use "snmpset" (OID 1.3.6.1.4.1.35098.2.0.0) to load a configuration and can use "snmpset" (OID 1.3.6.1.4.1.35098.2.1.0) to delete or load a configuration. However, only set and delete commands can be included in the command batch (which is OID 1.3.6.1.4.1.35098.2.1.0). Other commands are invalid and ignored. Note that clearing a dependent configuration is not allowed.

```

admin@XorPlus# set protocols snmp community private authorization read-write
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Using snmpset to load a filter configuration:

```

root@dev:~# snmpset -v 2c -c private IP .1.3.6.1.4.1.35098.2.0.0 s
"tftp:1.1.5.1:/pica8/acl.conf"
iso.3.6.1.4.1.35098.2.0.0 = STRING: "tftp:1.1.5.1:/pica8/acl.conf"

```

Using snmpset to delete a filter configuration:

```

root@dev:~# snmpset -v 2c -c private IP .1.3.6.1.4.1.35098.2.1.0 s
"tftp:1.1.5.1:/pica8/delete-acl.conf"
iso.3.6.1.4.1.35098.2.0.0 = STRING: "tftp:1.1.5.1:/pica8/delete-acl.conf"

```

Configuring the Syslog Log Level

Configuring the Syslog Level

Our system has five syslog levels. Listed in order of most fatal to least fatal, the levels are: Fatal, Error, Warning, Info, and Trace. By default, the system is set to the Warning level. You can, of course, change the log level.

In the example below, the system logs messages from Info, Warning, Error, and Fatal levels since the system syslog level is set to Info.

```

admin@XorPlus# set system log-level info
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
You can display the log messages on the console screen by entering the following command:
admin@XorPlus# exit
admin@XorPlus> syslog monitor on
If the switch's syslog level is Trace, the trace options of the modules should be turned
on, as illustrated below. You can also turn on the OSPF trace options for debugging.
admin@XorPlus# set protocols ospf4 traceoptions flag all disable false
admin@XorPlus# set system log-level trace
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# exit
admin@XorPlus> syslog monitor on

```

Configuring the SNMP Logging Facility

In accordance with the syslog standard, the logging facility can be configured as [0, 7].

```

admin@XorPlus# set system log-facility 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
Oct 17 15:22:42 XorPlus local0.warn : admin logged the switch
Oct 17 15:22:50 XorPlus local0.warn pica_sh: Tacacs send acct body send failed: wrote -1 of
127: Connection refused
admin@XorPlus# set system log-facility 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Oct 17 15:22:42 XorPlus local2.warn : admin logged the switch

Configuring the Syslog Disk and Syslog Host

Configuring the Syslog Host

After you configure the syslog server IP address, the log files will be sent to the syslog server.

```
admin@XorPlus# set system syslog server-ip 192.168.1.1 ?
Possible completions:
  <[Enter]>                Execute this command
  port                    Remote syslog server port
  protocol                Remote syslog server protocol
admin@XorPlus# set system syslog server-ip 192.168.1.1 protocol tcp
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```



Configuration of nodes "system syslog host", "system syslog port-number", and "system syslog port-protocol" have been deprecated in version 2.4.

Configuring Syslog for Local Storage

You can configure syslog messages to be stored in RAM or in a local SD card.

```
XorPlus# set system syslog local-file disk
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus#
XorPlus# set system syslog local-file ram
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus#
```

Displaying System Information

User can display a system's information, including fan, power supply unit, and serial number information.

Displaying the System Fan

```
admin@XorPlus>show system fan
Sensor Temperature:
Sensor 1 Temperature : 42 Centigrade
Sensor 2 Temperature : 39 Centigrade
Sensor 3 Temperature : 46 Centigrade
Sensor 4 Temperature : 33 Centigrade
Fan Status:
Fan 1 speed = 12529 RPM, PWM = 79
Fan 2 speed = 12413 RPM, PWM = 79
Fan 3 speed = 12300 RPM, PWM = 79
```

Displaying the System Power Supply Unit

```
admin@XorPlus> show system rpsu
RPSU 1:
TEMPERATURE_1 : N/A
RPSU 2:
TEMPERATURE_1 : 38.00 Centigrade
TEMPERATURE_2 : 40.00 Centigrade
FAN_SPEED : 10784.0 RPM
FAN_PWM : 60
```

Displaying the System Serial Number

```
admin@XorPlus> show system serial-number
MotherBoard Serial Number : QTFCXI2460009
RPSU 1 Serial Number : N/A
RPSU 2 Serial Number : 601G10103C370ZG
SFP te-1/1/49 :
Vendor Name : PICA8
Serial Number : 78613B10987
Module Type : SR/850nm
Cable Length : 80m
SFP te-1/1/50 :
Vendor Name : JESS-LINK
Serial Number : 12344D0001
Cable Length : 5m
SFP te-1/1/51 :
Vendor Name : DELTA
Serial Number : 084109000017
Module Type : SR/850nm
Cable Length : 80m
SFP te-1/1/52 :
Vendor Name : JESS-LINK
Serial Number : 12344D0002
Cable Length : 5m
```

Displaying Additional System Information

```
admin@XorPlus# run show system temperature
Temperature: 39 C /102F
```

```

admin@XorPlus#
admin@XorPlus# run show system uptime
01:21:33 up 50 min, load average: 0.04, 0.06, 0.07
admin@XorPlus#
admin@XorPlus# run show system cpu-usage
Cpu usage: 15%
admin@XorPlus#
admin@XorPlus# run show system date
Mon Jan 13 18:11:04 UTC 2014
admin@XorPlus#
admin@XorPlus# run show system memory-usage
total used free shared buffers cached
Mem: 515808 185468 330340 0 10320 68312
-/+ buffers/cache: 106836 408972
Swap: 0 0 0
admin@XorPlus#
admin@XorPlus# run show system name
admin@XorPlus
admin@XorPlus#
admin@XorPlus# run show system ntp-status
Please start the ntp server first!
admin@XorPlus#
admin@XorPlus# run show system os
Linux XorPlus 2.6.27 #1 Thu Feb 13 00:42:23 CST 2014 ppc GNU/Linux
admin@XorPlus# run show system processes brief
PID TTY STAT TIME COMMAND
1 ? Ss 0:01 init [2]
2 ? S< 0:00 [kthreadd]
3 ? S< 0:00 [ksoftirqd/0]
4 ? S< 0:00 [watchdog/0]
5 ? S< 0:02 [events/0]
6 ? S< 0:00 [khelper]
48 ? S< 0:00 [kblockd/0]
55 ? S< 0:00 [ata/0]
56 ? S< 0:00 [ata_aux]
58 ? S< 0:00 [kseriod]
99 ? S 0:00 [pdflush]
101 ? S< 0:00 [kswapd0]
147 ? S< 0:00 [aio/0]
156 ? S< 0:00 [nfsiod]
831 ? S< 0:00 [ftld]
853 ? S< 0:00 [rpciod/0]
857 ? S< 0:00 [kjournald]
2222 ? S 0:00 [pdflush]
2356 ? Ss 0:00 /usr/sbin/cron -L 0
2387 ? Ss 0:00 /usr/sbin/xinetd -pidfile /var/run/xinetd.pid -stayalive -inetd_compat
-inetd_ipv6
2501 ? S 0:03 pica_cardmgr
2503 ? S 0:59 pica_sif
2649 ? S 0:05 pica_lacp
2664 ? Ss 0:00 dhclient -pf /run/dhclient.eth0.pid -lf /var/lib/dhcp/dhclient.eth0.leases
eth0
2666 ? Sl 18:06 pica_lcmgr
2672 ? S 0:04 pica_login
3166 ? Sl 0:00 /usr/sbin/rsyslogd -c5
3457 ? S 0:35 pica_mstp
3462 ? S 0:02 xorp_policy
3464 ? Ss 1:03 /pica/bin/xorp_rtrmgr -d -L local0.info -P /var/run/xorp_rtrmgr.pid
3500 tty1 Ss+ 0:00 /sbin/getty 38400 tty1
3507 tty2 Ss+ 0:00 /sbin/getty 38400 tty2
3508 tty3 Ss+ 0:00 /sbin/getty 38400 tty3
3761 ttyS0 Ss+ 0:00 /sbin/getty -s -L ttyS0 115200 ansi
4050 ? S 0:57 ovs-vswitchd
4422 ? Ss 0:00 in.telnetd: 10.10.50.16
4423 pts/0 Ss 0:00 login -h 10.10.50.16 -p
4424 pts/0 S+ 0:00 -bash
4434 pts/0 S+ 0:03 /pica/bin/pica_sh
6451 ? Ss 0:00 in.telnetd: 10.10.50.18
6452 pts/1 Ss 0:00 login -h 10.10.50.18 -p
6460 pts/1 S+ 0:00 -bash

```



```

-rw-rw-r-- 1 root xorp 23595 Jul 7 22:26 /pica/config/pica.conf.03
admin@XorPlus# run show system users
admin pts/0 Jan 13 14:19 (10.10.50.16)
admin pts/1 Jan 13 15:03 (10.10.50.18)
admin@XorPlus#
admin@XorPlus# run show system core-dumps
total 0
admin@XorPlus#
admin@XorPlus# run show system connections
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address State User Inode
tcp 0 0 127.0.0.1:49152 0.0.0.0:* LISTEN 0 6787
tcp 0 0 127.0.0.1:60833 0.0.0.0:* LISTEN 0 5715
tcp 0 0 127.0.0.1:51714 0.0.0.0:* LISTEN 11 31043
tcp 0 0 127.0.0.1:42179 0.0.0.0:* LISTEN 0 6789
tcp 0 0 127.0.0.1:56484 0.0.0.0:* LISTEN 0 5711
tcp 0 0 127.0.0.1:51044 0.0.0.0:* LISTEN 0 5705
tcp 0 0 127.0.0.1:40421 0.0.0.0:* LISTEN 0 6764
tcp 0 0 127.0.0.1:56263 0.0.0.0:* LISTEN 0 6822
admin@XorPlus# run show system boot-messages
Copyright (c) 2009-2014 Pica8 Inc.
All rights reserved.
Up time: 18:19:41
revision: 2.6.27
Using MPC85xx CDS machine description
Memory CAM mapping: CAM0=256Mb, CAM1=256Mb, CAM2=0Mb residual: 0Mb
Linux version 2.6.27 (root@dev-16-new) (gcc version 4.2.2) #1 Thu Feb 13 00:42:23 CST 2014
Found legacy serial port 0 for /soc8541@e0000000/serial@4500
mem=e0004500, taddr=e0004500, irq=0, clk=330000000, speed=0
Found legacy serial port 1 for /soc8541@e0000000/serial@4600
mem=e0004600, taddr=e0004600, irq=0, clk=330000000, speed=0

```

Technical Support

Execute the diagnostic command *show tech_support* to send information to Pica8 Technical Support and receive a diagnostic report back.

Executing the Diagnostic Command

```

admin@XorPlus> show tech_support
Start.....
Item 1: Display system version finished!
Item 2: Display system interface finished!
Item 3: Display system configuration finished!
Item 4: Display system config files finished!
Item 5: Display system process finished!
Item 6: Display system fdb table finished!
Item 7: Display system fdb entries finished!
Item 8: Display system ospf neighbors finished!
Item 9: Display system ospf interfaces finished!
Item 10: Display system route table finished!
Item 11: Get error event from log!
Item 12: Display system hard-route table finished!
Item 13: Display system hard-route for host finished!
Item 14: Display system spanning tree interfaces finished!
Item 15: Display system spanning tree bridge finished!
Item 16: Display system vlans table finished!
Item 17: Display system vlan-interfaces finished!
Item 18: Display system core-dump finished!
Item 19: Display system uptime finished!
Item 20: Display system arp table finished!
The information has been stored in /tmp/XorPlus-201307052220-techSupport.log, please
forward to support@pica8.com
admin@XorPlus>

```

Flushing ARP and the Neighbor Table

You can manually flush the ARP entry and the IPv6 neighbor table.

Flushing the ARP Entry

```
admin@XorPlus> flush arp all
admin@XorPlus> flush arp ip-address 192.168.1.1
```

Rebooting the System

```
admin@XorPlus>request system reboot
U-Boot 1.3.0 (Apr 11 2011 - 10:41:10)
CPU: 8541, Version: 1.1, (0x80720011)
Core: E500, Version: 2.0, (0x80200020)
Clock Configuration:
CPU: 825 MHz, CCB: 330 MHz,
DDR: 165 MHz, LBC: 41 MHz
L1: D-cache 32 kB enabled
I-cache 32 kB enabled
I2C: ready
DRAM: Initializing
DDR: 512 MB
FLASH: 32 MB
L2 cache 256KB: enabled
Set ethaddr MAC address = 60:eb:69:d2:9c:d8
In: serial
Out: serial
Err: serial
Net: TSEC0
IDE: Bus 0: OK
Device 0: Model: TRANSCEND Firm: 20091130 Ser#: 20100723 C4130E83
Type: Hard Disk
Capacity: 1911.6 MB = 1.8 GB (3915072 x 512)
```

Displaying the Debugging Message

User can configure the debugging message in a current window.

Syslog Monitor On

```

admin@XorPlus> syslog monitor on
Nov 21 2000 22:27:39 XorPlus local0.warn : [SIF]Interface ge-1/1/3, changed state to up
Nov 21 2000 22:27:41 XorPlus local0.warn : root logged the switch
Nov 21 2000 22:41:18 XorPlus local0.info xinetd[1102]: START: telnet pid=7650
from=10.10.50.16
Nov 21 2000 22:41:23 XorPlus authpriv.debug login[7651]: pam_unix(login:account): account
admin has password changed in future
Nov 21 2000 22:41:26 XorPlus local0.warn : admin logged the switch
Nov 21 2000 22:55:58 XorPlus local0.info xinetd[1102]: START: telnet pid=8039
from=10.10.51.16
Nov 21 2000 22:56:01 XorPlus authpriv.debug login[8040]: pam_unix(login:account): account
root has password changed in future
Nov 21 2000 23:31:13 XorPlus local0.info xinetd[1102]: START: telnet pid=9028
from=10.10.50.16
Nov 21 2000 23:31:16 XorPlus authpriv.debug login[9029]: pam_unix(login:account): account
admin has password changed in future
Nov 21 2000 23:31:21 XorPlus local0.warn : admin logged the switch
admin@XorPlus>

```

System Management Command List

User can view the list of system management commands supported by the Pica8 PicOS software here.

IPv6 Management Support

SNMP over IPv6

User can walk the SNMP model via IPv6 address.

Telnet/SSH over IPv6

User can log in to the switch via IPv6 address. (Static IPv6 address for a management port can be configured.)

```

admin@XorPlus# set system management-ethernet eth0 ip-address ?
Possible completions:
<[Enter]> Execute this command
IPv4 IPv4 address or negotiated via DHCP, e.g. 192.168.1.2/24 or dhcp
IPv6 IPv6 address or negotiated via DHCP, e.g. fec0::10/64 or dhcp
admin@XorPlus# set system management-ethernet eth0 ip-gateway ?
Possible completions:
<[Enter]> Execute this command
IPv4 Configure the IPv4 gateway
IPv6 Configure the IPv6 gateway

```

Syslog over IPv6

User can send syslog via IPv6 address to remote server.

```
admin@XorPlus# set system syslog server-ip ?
Possible completions:
<IPv4> Remove syslog server
<IPv6> Remove syslog server for IPv6
```

NTP Client IPv6

User can synchronize the time via IPv6 address.

```
admin@XorPlus# set system ntp-server-ip ?
Possible completions:
<IPv4> Sync time with NTP server <IP>
<IPv6> Sync time with NTP server <IPv6>
```

DNS Client IPv6

User can configure an IPv6 dns server IP address.

```
admin@XorPlus# set system dns-server-ip ?
Possible completions:
<IPv4> DNS server IP address <IP>
<IPv6> DNS server IPv6 address <IP>
```

DHCPv6 Client

PicOS supports dhcpv6 function.

Radius over IPv6

User can configure an IPv6 address for the radius server.

```
admin@XorPlus# set system aaa radius authorization server-ip ?
Possible completions:
<IPv4> Radius authorization IPv4 server address
<IPv6> Radius authorization IPv6 server address
```

PoE Configuration

PoE Configuration



PoE commands are only supported on PoE enabled devices. Currently, of the devices on our Hardware compatibility list, only the AS4610-54P and the AS4610-30P are PoE enabled.

Show Help Information

```
root@XorPlus$ ./poe_tool -help
```

Show Version Information

```
root@XorPlus$ ./poe_tool -v
```

Key Words

Key words contain "-o, -p, -P". Meanings of these suffixes are as follows:

-o refers to operation code (key). It is a hexadecimal number, such as 0x19.

-p refers to port number. Port numbers range from 0 to 48. When p is 0, all ports are referred to.

-P refers to parameters. Three is the maximum number of parameters. Multiple parameters should be separated by semicolons. (e.g. p1;p2;p3).

PSE Function Enable

Key:0x00

The number of parameter is 1. The values range from 0 to 3.

0: disable pse

1: enable pse

2: force power

3: reserved

Configure enable pse for all ports:

```
root@XorPlus$ ./poe_tool -o 0x00 -p 0 -P 1
```

Configure force power for port 1:

```
root@XorPlus$ ./poe_tool -o 0x00 -p 1 -P 2
```

Configure fail returns:

Port xx PSE function enable/disable/set force power failed

Configure success returns

Port xx PSE function enable/disable/set force power success

If user selects parameter 3, nothing is returned.

Powerup Command

Key:0x01

The number of parameter is 1. The values range from 0 to 1.

0disallow powerup

1allow powerup

Configure disallow powerup for all ports:

```
root@XorPlus$ ./poe_tool -o 0x01 -p 0 -P 0
```

Configure allow powerup for port 32:

```
root@XorPlus$ ./poe_tool -o 0x01 -p 32 -P 1
```

Configure fail returns:

Port xx set powerup/powerdown fail

Configure success returns

Port xx set powerup/powerdown success

Logical Port Map Enable

Key:0x02

The number of parameter is 1. The values range from 0 to 1.

0disallow logical port map

1allow logical port map

Configure disallow logical port map:

```
root@XorPlus$ ./poe_tool -o 0x02 -P 0
```

Configure allow logical port map:

```
root@XorPlus$ ./poe_tool -o 0x02 -P 1
```

Configure fail returns:

Enable/Disable logical port map fail

Configure success returns

Enable/Disable logical port map success

Reset the Specified Port or Ports

Key:0x03

The number of parameter is 1. The values range from 0 to 1.

0: do not reset the port

1: reset the port

Configure reset port 1:

```
root@XorPlus$ ./poe_tool -o 0x03 -p 1 -P 1
```

Configure reset all ports:

```
root@XorPlus$ ./poe_tool -o 0x03 -p 0 -P 1
```

Configure fail returns:

Reset port xx fail

Configure success returns

Reset port xx success

If user selects parameter 0, nothing is returned.

Reset the Error Statistics on Specified Ports

Key:0x05

The number of parameter is 1. The values range from 0 to 1.

0: don't reset the port error statistics

1: reset the port error statistics

Reset the error statistics on port 1:

```
root@XorPlus$ ./poe_tool -o 0x05 -p 1 -P 1
```

Reset the error statistics on all ports:

```
root@XorPlus$ ./poe_tool -o 0x05 -p 0 -P 1
```

Configure fail returns:

Port xx statistics reset fail

Configure success returns

Port xx statistics reset success

If user selects parameter 0, nothing is returned.

Set the PSE Functionality on all Ports

Key:0x06

The number of parameter is 1. The values range from 0 to 3.

0: disable PSE functionality on all ports

1: enable PSE functionality on all ports

2: enable force power functionality on all ports

3: enable force power disconnect functionality on all ports

Configure enable PSE functionality on all ports:

```
root@XorPlus$ ./poe_tool -o 0x06 -P 1
```

Configure fail returns:

Disable PSE functionality/Enable PSE functionality/Enable force power functionality/Enable force power with disconnect functionality on all ports fail

Configure success returns

Disable PSE functionality/Enable PSE functionality/Enable force power functionality/Enable force power with disconnect functionality on all ports success

Set the Default Maximum Allocated Power Parameters in High Power Mode on all Ports in the PoE Subsystem

Key:0x07

The number of parameters is 1. The values range from 0 to 3.

0: 22.5W

1: 26.5W

2: 31.2W

3: 37 W

Configure default maximum. Allocated power is 26.5W in high power mode on all ports in the PoE subsystem.

```
root@XorPlus$ ./poe_tool -o 0x07 -P 1
```

Configure fail returns:

Set default maximum allocated power on all ports fail

Configure success returns

Set default maximum allocated power on all ports success

The Switch Port from Low Power Mode to High Power Mode

Key:0x08

The number of parameter is 1. The values range from 0 to 1.

0:switch the port from low power to high power

1:ignore

Configure all switch port from low power mode to high power mode.

```
root@XorPlus$ ./poe_tool -o 0x08 -p 0 -P 0
```

Configure fail returns:

set port xx to <high power mod> fail

Configure success returns

set port xx to <high power mod> success

If user selects parameter 1, nothing is returned.

Reset the Entire PoE Subsystem

Key:0x09

The number of parameter is 1. The values range from 0 to 1.

0: ignore

1: reset the PoE subsystem

Configure reset the entire PoE subsystem.

```
root@XorPlus$ ./poe_tool -o 0x09 -P 1
```

Configure fail returns:

PoE reset fail

Configure success returns

PoE reset success

Global System Parameters Configuration Command

Key:0x0a

The number of parameter is 3.

parameter 1: UVLO (Under Voltage Lock Out,the values should be bigger than 33,recommended to choose 40)

parameter 2: 0-1 (0:disable double detection 1: enable double detection)

parameter 3: OVLO (Over Voltage Locked-out,the values should be bigger than 57,Recommended to choose 60)

Configure UVLO 40, enable double detection, OVLO 60.

```
root@XorPlus$ ./poe_tool -o 0x0a -P 40:1:60
```

Configure fail returns:

Global system parameters configuration: operation fail

Configure success returns

Global system parameters configuration: operation success

Configure Detection Type on Specified Ports**Key: 0x10**

The number of parameter is 1. The values range from 0 to 6.

0: no detection

1: Legacy Capacitive Detection only

2: IEEE 802.3af 4-Point Detection only

3: IEEE 802.3af 4-Point followed by Legacy

4: IEEE 802.3af 2-Point Detection

5: IEEE 802.3af 2-Point followed by Legacy

6: Reserved

Configure no detection for all port:

```
root@XorPlus$ ./poe_tool -o 0x10 -p 0 -P 0
```

Configure fail returns:

Port xx set detection type to <detection type> fail

Configure success returns:

Port xx set detection type to <detection type> success

Configure the Classify Type on Specific Ports**Key: 0x11**

The number of parameters is 1. The values range from 0 to 1.

0: bypass classification

1: enable classification

Configure enable classification for all ports:

```
root@XorPlus$ ./poe_tool -o 0x11 -p 0 -P 1
```

Configure fail returns:

Port xx bypass/enable classification configuration fail

Configure success returns:

Port xx bypass/enable classification configuration success

Configure Auto Mode on Specific Ports

Key: 0x12

The number of parameter is 1. The values range from 0 to 1.

0: disable auto mode

1: enable auto mode

Configures auto mode on all ports:

```
root@XorPlus$ ./poe_tool -o 0x12 -p 0 -P 1
```

Configure fail returns:

Disable/enable port xx auto power-up mode fail

Configure success returns:

Disable/enable port xx auto power-up mode success

Configure Disconnect Type on Specific Ports

Key :0x13

The number of parameter is 1. The values range from 0 to 3.

0: None

1: AC Disconnect

2: DC Disconnect

3: DC Disconnect with delay

Configure DC Disconnect for all ports:

```
root@XorPlus$ ./poe_tool -o 0x13 -p 0 -P 2
```

Configure fail returns:

Port xx set disconnect type to <disconnect type> fail

Configure success returns:

Port xx set disconnect type to <disconnect type> success

Configure Power Threshold Type on Specified Ports

Key:0x15

The number of parameter is 1. The values range from 0 to 2.

0: None

1: Class Based

2: User defined

Configure Class Based for all ports:

```
root@XorPlus$ ./poe_tool -o 0x15 -p 0 -P 1
```

Configure fail returns:

Port xx set power threshold type to <power threshold type> fail

Configure success returns:

Port xx set power threshold type to <power threshold type> success

Configure Max Threshold on Specific Ports

Key:0x16

The number of parameter is 1. The values should not be greater than 30.

Configure max threshold 10 for port 1:

```
root@XorPlus$ ./poe_tool -o 0x16 -p 1 -P 10
```

Configure fail returns:

Port xx power threshold set to <max threshold> W fail

Configure success returns:

Port xx power threshold set to <max threshold> W success

Configure Power Management Mode

Key:0x17

The number of parameter is 1. The values range from 0 to 4.

0: None

1: Static Power Management with Port Priority

2: Dynamic Power Management with Port Priority

3: Static Power Management without Port Priority

4: Dynamic Power Management without Port Priority

Configure Static Power Management with Port Priority:

```
root@XorPlus$ ./poe_tool -o 0x17 -P 1
```

Configure fail returns:

Set power management mod to <power management mode> fail

Configure success returns:

Set power management mod to <power management mode> success

Configure Power Supply Details for Each Combination of the MPSS Pins

Key:0x18

The number of parameter is 3.

parameter 1: multi power source status pin combination,the values from 0 to 7.

parameter 2: total power. the values from 0 to 1600.

parameter 3: Guard band (recommend 10 percent of total power)

Configure power supply details for each combination of the mpss pins:

```
root@XorPlus$ ./poe_tool -o 0x18 -P 0:100:10
```

Configure fail returns:

Config power supply to PoE system fail

Configure success returns:

Config power supply to PoE system success

Configure Port Priority on Specific Ports

Key: 0x1a

The number of parameter is 1. The values range from 0 to 3.

0: low

1: medium

2: high

3: critical

Configure critical priority for port 1:

```
root@XorPlus$ ./poe_tool -o 0x1a -p 1 -P 3
```

Configure fail returns:

Port xx port priority configuration set to <port priority> fail

Configure success returns:

Port xx port priority configuration set to <port priority> success

Configures Port Power Up Mode

Key: 0x1c

The number of parameter is 1. The values from 0 to 3.

0: IEEE 802.3af mode

1: high inrush mode

2: pre-IEEE 802.3at

3: IEEE 802.3at mode

Configure:

```
root@XorPlus$ ./poe_tool -o 0x1c -p 0 -P 0
```

Configure fail returns:

Port xx set power up mod fail

Configure success returns:

Port xx set power up mod success

Configure the Per Port Power Disconnect Mask

Key: 0x1e

The number of parameter is 1. The values range from 0 to 255.

Configure power disconnect mask 17 for all ports

```
root@XorPlus$ ./poe_tool -o 0x1e -p 0 -P 17
```

Configure fail returns:

Port xx power disconnect mask configuration fail

Configure success returns:

Port xx power disconnect mask configuration success

Acquire Port Status

Key: 0x21

The number of parameter is 0.

Acquire port 1 status:

```
root@XorPlus$ ./poe_tool -o 0x21 -p 1
```

Return port status

Acquires the Total Power Allocated

Key: 0x23

The number of parameter is 0.

Acquires the total power allocated:

```
root@XorPlus$ ./poe_tool -o 0x23
```

Return:

Total power allocated to the ports is xx W. Power available is xx W.

NETCONF Configuration

NETCONF is a network configuration and management protocol based on XML.

NETCONF protocol uses XML for configuration data and protocol message encoding, using RPC and Client/Server mechanism to update, install or delete the relevant part of the device configuration or all the management information.



Now, we support <get>, <get-config> and <edit-config>

Enable NETCONF on switch:

```
admin@XorPlus# set protocols netconf
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

Delete NETCONF configuration on switch:

```
admin@XorPlus# delete protocols netconf
Deleting:
  netconf {
  }
OK
admin@XorPlus# commit
Commit OK.
Save done.
```

YANG is a data modeling language used to model configuration and state data manipulated by NETCONF.

You can find the YANG module file of different modules on your switch under the directory "/pica/bin/netconf/data-models".

```

admin@XorPlus$pwd
/pica/bin/netconf/data-models
admin@XorPlus$
admin@XorPlus$ls -lt *.yang
-rw-r--r-- 1 root xorp 2075 Jul 13 2016 arp.yang
-rw-r--r-- 1 root xorp 1143 Jul 13 2016 bfd.yang
-rw-r--r-- 1 root xorp 7767 Jul 13 2016 bgp.yang
-rw-r--r-- 1 root xorp 5612 Jul 13 2016 cos.yang
-rw-r--r-- 1 root xorp 2160 Jul 13 2016 dhcp.yang
-rw-r--r-- 1 root xorp 3030 Jul 13 2016 dot1x.yang
-rw-r--r-- 1 root xorp 13460 Jul 13 2016 firewall.yang
-rw-r--r-- 1 root xorp 16760 Jul 13 2016 ietf-inet-types.yang
-rw-r--r-- 1 root xorp 18034 Jul 13 2016 ietf-yang-types.yang
-rw-r--r-- 1 root xorp 1529 Jul 13 2016 igmp.yang
-rw-r--r-- 1 root xorp 2657 Jul 13 2016 igmpsnooping.yang
-rw-r--r-- 1 root xorp 41585 Jul 13 2016 interface.yang
-rw-r--r-- 1 root xorp 4991 Jul 13 2016 ipfix.yang
-rw-r--r-- 1 root xorp 1147 Jul 13 2016 lacp.yang
-rw-r--r-- 1 root xorp 3432 Jul 13 2016 lldp.yang
-rw-r--r-- 1 root xorp 781 Jul 13 2016 mfea.yang
-rw-r--r-- 1 root xorp 18188 Jul 13 2016 mstp.yang
-rw-r--r-- 1 root xorp 4673 Jul 13 2016 neighbour.yang
-rw-r--r-- 1 root xorp 7835 Jul 13 2016 ospfv2.yang
-rw-r--r-- 1 root xorp 5512 Jul 13 2016 ospfv3.yang
-rw-r--r-- 1 root xorp 3408 Jul 13 2016 ovsdb.yang
-rw-r--r-- 1 root xorp 4202 Jul 13 2016 pim.yang
-rw-r--r-- 1 root xorp 8203 Jul 13 2016 policy.yang
-rw-r--r-- 1 root xorp 3954 Jul 13 2016 rip.yang
-rw-r--r-- 1 root xorp 3031 Jul 13 2016 ripng.yang
-rw-r--r-- 1 root xorp 4624 Jul 13 2016 sflow.yang
-rw-r--r-- 1 root xorp 1207 Jul 13 2016 snmp.yang
-rw-r--r-- 1 root xorp 2666 Jul 13 2016 static-routes.yang
-rw-r--r-- 1 root xorp 1670 Jul 13 2016 stm.yang
-rw-r--r-- 1 root xorp 1666 Jul 13 2016 udd.yang
-rw-r--r-- 1 root xorp 2747 Jul 13 2016 vlan-interface.yang
-rw-r--r-- 1 root xorp 6063 Jul 13 2016 vlans.yang
-rw-r--r-- 1 root xorp 1747 Jul 13 2016 vrrp.yang
-rw-r--r-- 1 root xorp 6423 Jul 13 2016 vxlans.yang
-rw-r--r-- 1 root xorp 4186 Jul 13 2016 xovs.yang

```

Example of VLAN configuration via NETCONF use <edit-config>:

Step 1: Create an XML file according the vlan.yang for RPC request to create VLAN136:

```

<vlans xmlns="http://pica8.com/xorpplus/vlans">
  <id>136</id>
  <description/>
  <vlan-name>default</vlan-name>
  <l3-interface>vlan136</l3-interface>
</vlan-id>
</vlans>

```

Step 2: Display the configuration on switch after the client sending an RPC request.

```
The configuration has been changed by user root
DELTAS:
  vlans {
    vlan-id 136 {
      description: ""
      vlan-name: "default"
      l3-interface: "vlan136"
    }
  }
admin@XorPlus# show | display set
set protocols netconf
set vlans vlan-id 136 l3-interface "vlan136"
```



Now, we only support get the system's version information and vxlan information via NETCONF <get> function.

Example of get the system's version information via NETCONF use <get>:

Display the RPC reply after the client sending an RPC request.

```
<version xmlns="http://pica8.com/xorppplus/version">
  <mac_address>48:0f:cf:af:70:3b</mac_address>
  <hardware_mode>HP5712</hardware_mode>
  <system_version>2.8.0/aeec598</system_version>
  <system_released_date>10/13/2016</system_released_date>
  <L2_L3_version>2.8.0/aeec598</L2_L3_version>
  <L2_L3_released_date>10/13/2016</L2_L3_released_date>
</version>
```

NETCONF client

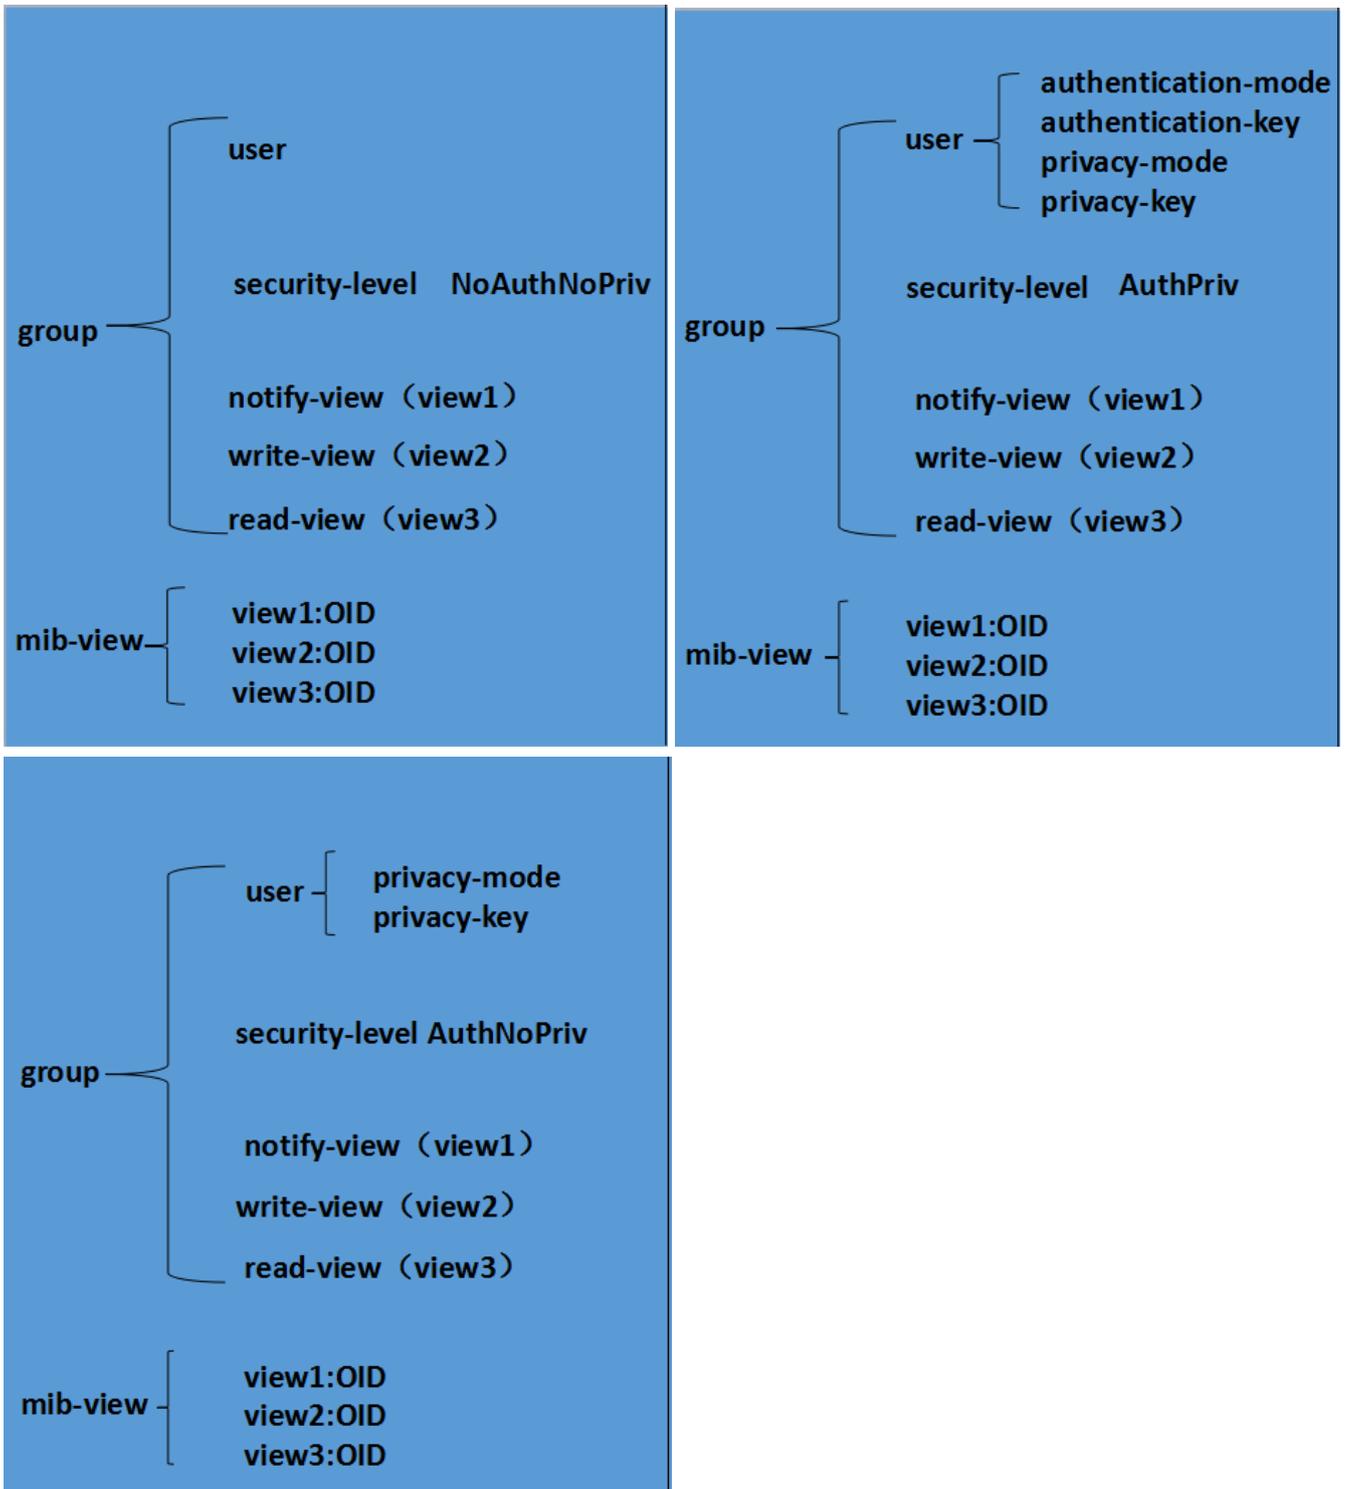
About NETCONF client, you can use ncclient which is python lib now.



If you use ncclient, you must modify the rpc.py : add two lines codes to work with pica8 switch.
 Edit the `rpc.py` file to contain the followings before the statement 'self._session.send(req)':
`req = req.replace('nc:', '')`
`req = req.replace(':', 'nc:')`

Configuring SNMPv3

Here is the configuration relation diagram in SNMPv3. A user can be added to a group or not as you need. Once a user joined a group, one or more kinds of views (notify-view, write-view, read-view) must be configured. Besides, configurations on user change with the security-level of the group as below 3 diagrams. Note notify-view, write-view and read-view are optional but you have to choose at least one view. In the below diagram, view1, view2, view3 can be the same or different.



Configuring Basic Information

Contact and location information can be configured as below which is the same as SNMPv2.

```
admin@XorPlus# set protocols snmp contact support@pica8.com
admin@XorPlus# set protocols snmp location beijing
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring trap-group

By default, trap messages are sent in the form of SNMPv2. But you can change it to SNMPv3 as below and designate NMS to which trap messages are sent. Note that in SNMPv3, security-name is user while in SNMPv2, security-name is community.

```
admin@XorPlus# set protocols snmp trap-group version v3
admin@XorPlus# set protocols snmp trap-group targets 10.10.51.42 security-name user1
admin@XorPlus# commit
Commit OK.
Save done.
```

Setting Up a User

By default, SNMPv3 is enabled. And you should set up a usm-user first before you configure other functions of SNMPv3. Besides, adding a user to a group is optional. When a user is added to a group, the needed views should be configured. If you create a user without adding to any groups, you can configure as below. However, under the below circumstance, all the OIDs can be visited by its NMS which can read but can't write and be notified.

```
admin@XorPlus# set protocols snmp v3 usm-user user1
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring Mib-view

If you want to improve security, the user needs to be added to a group. Because in this way, you can configure a read-view, write-view or notify-view (you can choose only one kind or more as you need) which defines the authority of a NMS. Before configuring a read-view (write-view or notify-view), please set up a mib-view which is used as a view of the group. Here are the configurations. As for a mib-view, you can include or exclude some subtrees and can also configure mask for them.

```
admin@XorPlus# set protocols snmp v3 usm-user user2 group group1
admin@XorPlus# set protocols snmp v3 mib-view view1 subtree 1.3.6.1.2.1 type included
admin@XorPlus# set protocols snmp v3 mib-view view1 subtree 1.3.6.1.2.1.6.13 type excluded

admin@XorPlus# set protocols snmp v3 mib-view view1 subtree 1.3.6.1.2.1 mask fc
admin@XorPlus# set protocols snmp v3 mib-view view1 subtree 1.3.6.1.2.1.6.13 mask ff
admin@XorPlus# set protocols snmp v3 group group1 read-view view1
admin@XorPlus# set protocols snmp v3 group group1 write-view view1
admin@XorPlus# set protocols snmp v3 group group1 notify-view view1
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring Security-level

You can improve security better by configuring security-level for the group. And the default setting of security-level is NoAuthNoPriv. You can change it to AuthNoPriv or AuthPriv. But please remember to configure authentication-mode, authentication-key, privacy-mode, privacy-key for the user. Configuration are as below.

```
admin@XorPlus# set protocols snmp v3 group group1 security-level AuthPriv
admin@XorPlus# set protocols snmp v3 usm-user user1 authentication-mode md5
admin@XorPlus# set protocols snmp v3 usm-user user1 authentication-key authnkey
admin@XorPlus# set protocols snmp v3 usm-user user1 privacy-mode des
admin@XorPlus# set protocols snmp v3 usm-user user1 privacy-key privykey
admin@XorPlus# commit
Commit OK.
Save done.
```

NMS Visits Switch by user

NMS reads OID tree 1.3.6.1.2.1.6.13 as below. user1 is the user's name. AuthPriv is the security-level of group. 10.10.51.155 is the IP of the switch.

```
pica8@pica8:~$snmpwalk -u user1 -l AuthPriv -A sha -a u1111key -X des -x u1111key
10.10.51.155 1.3.6.1.2.1.6.13
```

Ethernet Ports Configuration

- Physical Ethernet Port Configuration
- VLAN Port Configuration
- 40G Changes to 4*10G in L2/L3
 - 40G Changes to 4*10G in L2/L3 mode on P-5101
 - 40G Changes to 4*10G in L2/L3 mode on P-5401
 - 40G Changes to 4*10G in L2/L3 mode on accton_as6701_32x
 - 40G Changes to 4*10G in L2/L3 mode on P-3922
 - 40G Changes to 4*10G in L2/L3 mode on P-3920
 - 40G Changes to 4*10G in L2/L3 mode on accton_as5712_54x/HP5712
 - 40G Changes to 4*10G in L2/L3 mode on P-3930
 - 40G Changes to 4*10G in L2/L3 mode on Niagara2632XL
 - 40G Changes to 4*10G in L2/L3 mode on Niagara2948_6XL
 - 40G Changes to 4*10G in L2/L3 mode on as5812_54t
 - 40G Changes to 4*10G in L2/L3 mode on arctica4806xp
- LLDP Configuration (Link Layer Discovery Protocol)
- Voice Vlan with LLDP Configuration Guide
- Voice Vlan with LLDP Compliance CDP Configuration Guide
- Voice Vlan Configuration Guide
- LLDP MED Configuration
- Voice Vlan with LLDP MED Configuration Guide
- Static Link Aggregation (LAG) Configuration
- LAG Hashing Configuration
 - LAG Hashing Configuration and Example
 - LAG Hash Mapping
 - Resilient LAG Hashing Configuration and Example
- Link Aggregation Control Protocol (LACP) Configuration
- Storm Control in Ethernet Port Configuration
- UDLD Configuration
- Port Security Configuration
- 802.1x Configuration
- Buffer management

Physical Ethernet Port Configuration

You can enable (or disable) the Ethernet port, configure the Ethernet port's MTU, rate-limit, flow control, and change the qe-interface mode.

Shutting Down the Ethernet Port

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the MTU and Rate-limit

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 rate-limiting egress kilobits 10000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 mtu 1200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Enabling Port Flow Control

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options flow-control true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Split 40GE Ports into 4x10GE Ports

You can split 40GE ports to multiple 10GE ports using a split cable.



This is done by enabling the SFP mode on the relevant port. You can find details of this command in the command reference guide.

```
admin@XorPlus# set interface qe-interface-mode SFP
admin@XorPlus# com
Merging the configuration.
Commit OK.
Save done.
Qe interface mode changes, please reboot system to make it effect!
admin@XorPlus#
admin@XorPlus# set interface qe-interface-mode QSFP
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
Qe interface mode changes, please reboot system to make it effect!
admin@XorPlus#
```

Configuring Port Speed

```

admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1 detail
Physical interface: ge-1/1/1, Enabled, Physical link is Up
Interface index: 1
Link-level type: Ethernet, MTU: 1514, Speed: 1Gb/s, Duplex: Full
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit ingress:0, egress:0
Current address: c8:0a:a9:04:49:19, Hardware address: c8:0a:a9:04:49:19
Traffic statistics:
Input Packets.....35748
Output Packets.....35143881241
Input Octets.....3923150
Output Octets.....2266956387852
MAC statistics:
Multicast packets RX and TX.....199565932
Broadcast packets RX and TX.....4968094
Undersize packets RX and TX.....0
Fragments packets RX and TX.....0
Packets RX and TX 64 Octets.....35088774487
Packets RX and TX 65-127 Octets.....27771
Packets RX and TX 128-255 Octets.....2574126
Packets RX and TX 256-511 Octets.....52540605
Packets RX and TX 512-1023 Octets.....0
Packets RX and TX 1024-1518 Octets.....0
admin@XorPlus# run clear interface statistics all

```

VLAN Port Configuration

VLAN tagging (IEEE 802.1Q) is a networking standard that defines the VLAN. You can configure a port as either a trunk or an access port. With the native VLAN ID, you can add the port (in trunk mode) to more than one VLAN. Access ports belong to native VLANs, while trunk ports belong to more than one VLAN, including the native VLAN.

Configure the Access/Trunk Mode

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

Configure the Native VLAN ID

The native VLAN ID is the ID of the default VLAN (usually vlan-id 1) in which the port belongs. Every port should be included in at least one VLAN.

```

admin@XorPlus# set vlans vlan-id 5
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 5
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show vlans vlan-id 5
VLAN ID: 5
VLAN Name: default
Description:
vlan-interface:
Number of member ports: 1
Tagged port: None
Untagged port: ge-1/1/1,
admin@XorPlus#

```

Add a Port to a VLAN with Tagged Packet

```

admin@XorPlus# set vlans vlan-id 5
admin@XorPlus# set vlans vlan-id 6
admin@XorPlus# set vlans vlan-id 7
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 5
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 6
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 7
admin@XorPlus# commit
Commit OK.
Save done.

admin@XorPlus# run show vlans
VlanID  Tag          Interfaces
-----  -
1      untagged  ge-1/1/1, ge-1/1/2, ge-1/1/3, ge-1/1/4, ge-1/1/5
          ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9, ge-1/1/10
          ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14, ge-1/1/15
          ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19, ge-1/1/20
          ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24, te-1/1/25
          te-1/1/26, te-1/1/27, te-1/1/28, te-1/1/29, te-1/1/30
5      tagged
5      untagged
6      tagged    ge-1/1/2
6      untagged
6      tagged    ge-1/1/2
7      untagged
7      tagged    ge-1/1/2

```

Add a Port to a VLAN with Untagged Packet

```

admin@XorPlus# set vlans vlan-id 5
admin@XorPlus# set vlans vlan-id 6
admin@XorPlus# set vlans vlan-id 7
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 5 untagged
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 6 untagged
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 7 untagged
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show vlans
VlanID   Tag           Interfaces
-----
1        untagged      ge-1/1/1, ge-1/1/2, ge-1/1/3, ge-1/1/4, ge-1/1/5
          ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9, ge-1/1/10
          ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14, ge-1/1/15
          ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19, ge-1/1/20
          ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24, te-1/1/25
          te-1/1/26, te-1/1/27, te-1/1/28, te-1/1/29, te-1/1/30
          tagged
5        untagged      ge-1/1/2
          tagged
6        untagged      ge-1/1/2
          tagged
7        untagged      ge-1/1/2
          tagged

```

Creating a VLAN within the VLAN Range

You can create VLANs within the VLAN range, and then add ports to these VLANs.

```

admin@XorPlus# set vlans vlan-id 2-4094
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 1-4094
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 1-4094
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 1-4094
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

VLAN Configuration Example

In the following topology, the VLANs are configured for each switch.

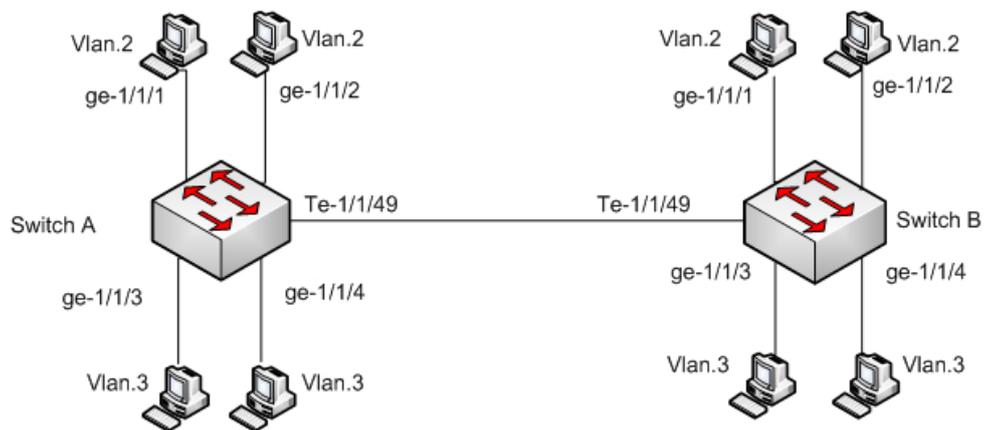


Figure 4-1. VLAN configuration.

Configuring Switch A

For Switch A, you should configure ge-1/1/1 through ge-1/1/4 as access ports and te-1/1/49 as the trunk port because the 10Gbit link will trunk the traffic of VLAN-2 and VLAN-3.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 3
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 2
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 3
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show vlans
VlanID   Tag                Interfaces
-----
1 tagged
  untagged    ge-1/1/5, ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9,
                ge-1/1/10, ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14,
                ge-1/1/15, ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19,
                ge-1/1/20, ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24,
                ge-1/1/25, ge-1/1/26, ge-1/1/27, ge-1/1/28, ge-1/1/29,
                ge-1/1/30, ge-1/1/31, ge-1/1/32, ge-1/1/33, ge-1/1/34,
                ge-1/1/35, ge-1/1/36, ge-1/1/37, ge-1/1/38, ge-1/1/39,
                ge-1/1/40, ge-1/1/41, ge-1/1/42, ge-1/1/43, ge-1/1/44,
                ge-1/1/45, ge-1/1/46, ge-1/1/47, ge-1/1/48, te-1/1/49,
                te-1/1/50, te-1/1/51, te-1/1/52,
2 tagged    te-1/1/49,
  untagged    ge-1/1/1, ge-1/1/2,
3 tagged    te-1/1/49,
  untagged    ge-1/1/3, ge-1/1/4,
admin@XorPlus#

```

Configuring Switch B

For Switch B, configure ge-1/1/1 through ge-1/1/4 as access ports and te-1/1/49 as the trunk port because the 10 Gbit link will trunk the traffic of VLAN-2 and VLAN-3.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching port-mode
access
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 3
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 2
admin@XorPlus#set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show vlans
VlanID      Tag                Interfaces
-----
1 tagged
  untagged      ge-1/1/5, ge-1/1/6, ge-1/1/7, ge-1/1/8, ge-1/1/9,
                 ge-1/1/10, ge-1/1/11, ge-1/1/12, ge-1/1/13, ge-1/1/14,
                 ge-1/1/15, ge-1/1/16, ge-1/1/17, ge-1/1/18, ge-1/1/19,
                 ge-1/1/20, ge-1/1/21, ge-1/1/22, ge-1/1/23, ge-1/1/24,
                 ge-1/1/25, ge-1/1/26, ge-1/1/27, ge-1/1/28, ge-1/1/29,
                 ge-1/1/30, ge-1/1/31, ge-1/1/32, ge-1/1/33, ge-1/1/34,
                 ge-1/1/35, ge-1/1/36, ge-1/1/37, ge-1/1/38, ge-1/1/39,
                 ge-1/1/40, ge-1/1/41, ge-1/1/42, ge-1/1/43, ge-1/1/44,
                 ge-1/1/45, ge-1/1/46, ge-1/1/47, ge-1/1/48, te-1/1/49,
                 te-1/1/50, te-1/1/51, te-1/1/52,
2 tagged      te-1/1/49,
  untagged      ge-1/1/1, ge-1/1/2,
3 tagged      te-1/1/49,
  untagged      ge-1/1/3, ge-1/1/4,

```

40G Changes to 4*10G in L2/L3

In L2/L3 mode, P-5401 ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: All 32 ports work in 40G QSFP mode.
2. SFP: Ports 1-12, and 17-28 work in 4*10G, ports 13-16, and 29-32 work in 40G.
3. SFP- 64: Ports 1-8, and 17-24 work in 4*10G, ports 9-16, and 25-32 work in 40G.

In L2/L3 mode, P-5101 ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: Ports 1-40 work in 10G, ports 41-48 work in 40G.

2. SFP: Ports 1-40 work in 10G, ports 41-48 work in 4*10G.

In L2/L3 mode, as6701_32x ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: All 32 ports work in 40G.
2. SFP: Ports 5-16, and 21-32 work in 4*10G, ports 1-4, and 17-20 work in 40G.

In L2/L3 mode, P-3922 ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: Ports 1-48 work in 10G, ports 48-52 work in 40G.
2. SFP: Ports 1-48 work in 10G, ports 48-52 work in 4*10G.

In L2/L3 mode, P-3920 ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: Ports 1-48 work in 10G, ports 48-52 work in 40G.
2. SFP: Ports 1-48 work in 10G, ports 48-52 work in 4*10G.

In L2/L3 mode, as5712_54x ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: Ports 1-48 work in 10G, ports 48-54 work in 40G.
2. SFP: Ports 1-48 work in 10G, ports 48-54 work in 4*10G.

In L2/L3 mode, Niagara2632XL ports can be configured to one of the following settings. Default is QSFP mode.

1. QSFP: All 32 ports work in 40G QSFP mode.
2. SFP: Ports 1-12, and 17-28 work in 4*10G, ports 13-16, and 29-32 work in 40G.
3. SFP- 64: Ports 1-8, and 17-24 work in 4*10G, ports 9-16, and 25-32 work in 40G.

- 40G Changes to 4*10G in L2/L3 mode on P-5101
- 40G Changes to 4*10G in L2/L3 mode on P-5401
- 40G Changes to 4*10G in L2/L3 mode on accton_as6701_32x
- 40G Changes to 4*10G in L2/L3 mode on P-3922
- 40G Changes to 4*10G in L2/L3 mode on P-3920
- 40G Changes to 4*10G in L2/L3 mode on accton_as5712_54x/HP5712
- 40G Changes to 4*10G in L2/L3 mode on P-3930
- 40G Changes to 4*10G in L2/L3 mode on Niagara2632XL
- 40G Changes to 4*10G in L2/L3 mode on Niagara2948_6XL
- 40G Changes to 4*10G in L2/L3 mode on as5812_54t

40G Changes to 4*10G in L2/L3 mode on P-5101

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP (8 x 40G+40*10G)

When ports are in QSFP mode, the mapping between physical port and the associated port/interface name is in the following table.

Physical Port number	L2/L3 port/interface name
1	te-1/1/1
2	te-1/1/2
3	te-1/1/3
4	te-1/1/4
5	te-1/1/5
6	te-1/1/6
7	te-1/1/7
8	te-1/1/8
9	te-1/1/9
10	te-1/1/10
11	te-1/1/11
12	te-1/1/12
13	te-1/1/13
14	te-1/1/14
15	te-1/1/15
16	te-1/1/16

17	te-1/1/17
18	te-1/1/18
19	te-1/1/19
20	te-1/1/20
21	te-1/1/21
22	te-1/1/22
23	te-1/1/23
24	te-1/1/24
25	te-1/1/25
26	te-1/1/26
27	te-1/1/27
28	te-1/1/28
29	te-1/1/29
30	te-1/1/30
31	te-1/1/31
32	te-1/1/32
33	te-1/1/33
34	te-1/1/34
35	te-1/1/35
36	te-1/1/36
37	te-1/1/37
38	te-1/1/38
39	te-1/1/39
40	te-1/1/40
41	qe -1/1/41
42	qe -1/1/42
43	qe -1/1/43
44	qe -1/1/44
45	qe -1/1/45
46	qe -1/1/46

47	qe -1/1/47
48	qe -1/1/48

SFP (72x 10G)

When ports are in SFP mode, the mapping between physical port and the associated port/interface name is in the following table.

Physical Port number	L2/L3 port/interface name
1	te-1/1/1
2	te-1/1/2
3	te-1/1/3
4	te-1/1/4
5	te-1/1/5
6	te-1/1/6
7	te-1/1/7
8	te-1/1/8
9	te-1/1/9
10	te-1/1/10
11	te-1/1/11
12	te-1/1/12
13	te-1/1/13
14	te-1/1/14
15	te-1/1/15
16	te-1/1/16
17	te-1/1/17
18	te-1/1/18
19	te-1/1/19
20	te-1/1/20
21	te-1/1/21
22	te-1/1/22
23	te-1/1/23

24	te-1/1/24
25	te-1/1/25
26	te-1/1/26
27	te-1/1/27
28	te-1/1/28
29	te-1/1/29
30	te-1/1/30
31	te-1/1/31
32	te-1/1/32
33	te-1/1/33
34	te-1/1/34
35	te-1/1/35
36	te-1/1/36
37	te-1/1/37
38	te-1/1/38
39	te-1/1/39
40	te-1/1/40
41	4 x 10G
	te -1/1/41
	te -1/1/42
	te -1/1/43
	te -1/1/44
42	4 x 10G
	te -1/1/45
	te -1/1/46
	te -1/1/47
	te -1/1/48
43	4 x 10G
	te -1/1/49
	te -1/1/50

	te -1/1/51
	te -1/1/52
44	4 x 10G
	te -1/1/53
	te -1/1/54
	te -1/1/55
	te -1/1/56
45	4 x 10G
	te -1/1/57
	te -1/1/58
	te -1/1/59
	te -1/1/60
46	4 x 10G
	te -1/1/61
	te -1/1/62
	te -1/1/63
	te -1/1/64
47	4 x 10G
	te -1/1/65
	te -1/1/66
	te -1/1/67
	te -1/1/68
48	4 x 10G
	te -1/1/69
	te -1/1/70
	te -1/1/71
	te -1/1/72

**Note:**

P-5101 does not support SFP-64 mode.

40G Changes to 4*10G in L2/L3 mode on P-5401

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/SFP-64/QSFP.
admin@XorPlus# commit
```

After setting ports to different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP (32 x 40G)

When ports are in QSFP mode, the mapping between physical ports, interface name, and interface support speed is in the following table.

Physical Port number	Interface name	Interface support speed
1	qe-1/1/1	40Gb/s
2	qe-1/1/2	40Gb/s
3	qe-1/1/3	40Gb/s
4	qe-1/1/4	40Gb/s
5	qe-1/1/5	40Gb/s
6	qe-1/1/6	40Gb/s
7	qe-1/1/7	40Gb/s
8	qe-1/1/8	40Gb/s
9	qe-1/1/9	40Gb/s
10	qe-1/1/10	40Gb/s
11	qe-1/1/11	40Gb/s
12	qe-1/1/12	40Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s

17	qe-1/1/17	40Gb/s
18	qe-1/1/18	40Gb/s
19	qe-1/1/19	40Gb/s
20	qe-1/1/20	40Gb/s
21	qe-1/1/21	40Gb/s
22	qe-1/1/22	40Gb/s
23	qe-1/1/23	40Gb/s
24	qe-1/1/24	40Gb/s
25	qe-1/1/25	40Gb/s
26	qe-1/1/26	40Gb/s
27	qe-1/1/27	40Gb/s
28	qe-1/1/28	40Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

SFP-64 (16 x 40G + 64 x 10G)

When ports are in SFP-64 mode, the mapping between physical ports, interface name, and interface support speed is in the following table.

Physical Port number	Interface name	Interface support speed
1	4 x 10G	
	te-1/1/1	10Gb/s
	te-1/1/2	10Gb/s
	te-1/1/3	10Gb/s
	te-1/1/4	10Gb/s
2	4 x 10G	
	te-1/1/5	10Gb/s
	te-1/1/6	10Gb/s
	te-1/1/7	10Gb/s

	te-1/1/8	10Gb/s
3	4 x 10G	
	te-1/1/9	10Gb/s
	te-1/1/10	10Gb/s
	te-1/1/11	10Gb/s
	te-1/1/12	10Gb/s
4	4 x 10G	
	te-1/1/13	10Gb/s
	te-1/1/14	10Gb/s
	te-1/1/15	10Gb/s
	te-1/1/16	10Gb/s
5	4 x 10G	
	te-1/1/17	10Gb/s
	te-1/1/18	10Gb/s
	te-1/1/19	10Gb/s
	te-1/1/20	10Gb/s
6	4 x 10G	
	te-1/1/21	10Gb/s
	te-1/1/22	10Gb/s
	te-1/1/23	10Gb/s
	te-1/1/24	10Gb/s
7	4 x 10G	
	te-1/1/25	10Gb/s
	te-1/1/26	10Gb/s
	te-1/1/27	10Gb/s
	te-1/1/28	10Gb/s
8	4 x 10G	
	te-1/1/29	10Gb/s
	te-1/1/30	10Gb/s
	te-1/1/31	10Gb/s

	te-1/1/32	10Gb/s
9	qe-1/1/9	40Gb/s
10	qe-1/1/10	40Gb/s
11	qe-1/1/11	40Gb/s
12	qe-1/1/12	40Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	4 x 10G	
	te-1/1/33	10Gb/s
	te-1/1/34	10Gb/s
	te-1/1/35	10Gb/s
	te-1/1/36	10Gb/s
18	4 x 10G	
	te-1/1/37	10Gb/s
	te-1/1/38	10Gb/s
	te-1/1/39	10Gb/s
	te-1/1/40	10Gb/s
19	4 x 10G	
	te-1/1/41	10Gb/s
	te-1/1/42	10Gb/s
	te-1/1/43	10Gb/s
	te-1/1/44	10Gb/s
20	4 x 10G	
	te-1/1/45	10Gb/s
	te-1/1/46	10Gb/s
	te-1/1/47	10Gb/s
	te-1/1/48	10Gb/s
21	4 x 10G	

	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
22	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
23	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
24	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
25	qe-1/1/25	40Gb/s
26	qe-1/1/26	40Gb/s
27	qe-1/1/27	40Gb/s
28	qe-1/1/28	40Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

SFP (8 x 40G + 96 x 10G)

When ports are in SFP mode, the mapping between physical port number, interfaces name, and interface support speed is in the following table.

Physical Port number	Interface name	Interface support speed
1	4 x 10G	
	te-1/1/1	10Gb/s
	te-1/1/2	10Gb/s
	te-1/1/3	10Gb/s
	te-1/1/4	10Gb/s
2	4 x 10G	
	te-1/1/5	10Gb/s
	te-1/1/6	10Gb/s
	te-1/1/7	10Gb/s
	te-1/1/8	10Gb/s
3	4 x 10G	
	te-1/1/9	10Gb/s
	te-1/1/10	10Gb/s
	te-1/1/11	10Gb/s
	te-1/1/12	10Gb/s
4	4 x 10G	
	te-1/1/13	10Gb/s
	te-1/1/14	10Gb/s
	te-1/1/15	10Gb/s
	te-1/1/16	10Gb/s
5	4 x 10G	
	te-1/1/17	10Gb/s
	te-1/1/18	10Gb/s
	te-1/1/19	10Gb/s
	te-1/1/20	10Gb/s

6	4 x 10G	
	te-1/1/21	10Gb/s
	te-1/1/22	10Gb/s
	te-1/1/23	10Gb/s
	te-1/1/24	10Gb/s
7	4 x 10G	
	te-1/1/25	10Gb/s
	te-1/1/26	10Gb/s
	te-1/1/27	10Gb/s
	te-1/1/28	10Gb/s
8	4 x 10G	
	te-1/1/29	10Gb/s
	te-1/1/30	10Gb/s
	te-1/1/31	10Gb/s
	te-1/1/32	10Gb/s
9	4 x 10G	
	te-1/1/33	10Gb/s
	te-1/1/34	10Gb/s
	te-1/1/35	10Gb/s
	te-1/1/36	10Gb/s
10	4 x 10G	
	te-1/1/37	10Gb/s
	te-1/1/38	10Gb/s
	te-1/1/39	10Gb/s
	te-1/1/40	10Gb/s
11	4 x 10G	
	te-1/1/41	10Gb/s
	te-1/1/42	10Gb/s
	te-1/1/43	10Gb/s
	te-1/1/44	10Gb/s

12	4 x 10G	
	te-1/1/45	10Gb/s
	te-1/1/46	10Gb/s
	te-1/1/47	10Gb/s
	te-1/1/48	10Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
18	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
19	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
20	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
21	4 x 10G	

	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
22	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s
23	4 x 10G	
	te-1/1/73	10Gb/s
	te-1/1/74	10Gb/s
	te-1/1/75	10Gb/s
	te-1/1/76	10Gb/s
24	4 x 10G	
	te-1/1/77	10Gb/s
	te-1/1/78	10Gb/s
	te-1/1/79	10Gb/s
	te-1/1/80	10Gb/s
25	4 x 10G	
	te-1/1/81	10Gb/s
	te-1/1/82	10Gb/s
	te-1/1/83	10Gb/s
	te-1/1/84	10Gb/s
26	4 x 10G	
	te-1/1/85	10Gb/s
	te-1/1/86	10Gb/s
	te-1/1/87	10Gb/s
	te-1/1/88	10Gb/s
27	4 x 10G	

	te-1/1/89	10Gb/s
	te-1/1/90	10Gb/s
	te-1/1/91	10Gb/s
	te-1/1/92	10Gb/s
28	4 x 10G	
	te-1/1/93	10Gb/s
	te-1/1/94	10Gb/s
	te-1/1/95	10Gb/s
	te-1/1/96	10Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

40G Changes to 4*10G in L2/L3 mode on accton_as6701_32x

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP (32 x 40G)

When ports are in QSFP mode, the mapping between physical ports and the logical ports/interfaces name is in the following table.

Physical Port number	L2/L3 port/interface name	interface support speed
1	qe-1/1/1	40Gb/s
2	qe-1/1/2	40Gb/s

3	qe-1/1/3	40Gb/s
4	qe-1/1/4	40Gb/s
5	qe-1/1/5	40Gb/s
6	qe-1/1/6	40Gb/s
7	qe-1/1/7	40Gb/s
8	qe-1/1/8	40Gb/s
9	qe-1/1/9	40Gb/s
10	qe-1/1/10	40Gb/s
11	qe-1/1/11	40Gb/s
12	qe-1/1/12	40Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	qe-1/1/17	40Gb/s
18	qe-1/1/18	40Gb/s
19	qe-1/1/19	40Gb/s
20	qe-1/1/20	40Gb/s
21	qe-1/1/21	40Gb/s
22	qe-1/1/22	40Gb/s
23	qe-1/1/23	40Gb/s
24	qe-1/1/24	40Gb/s
25	qe-1/1/25	40Gb/s
26	qe-1/1/26	40Gb/s
27	qe-1/1/27	40Gb/s
28	qe-1/1/28	40Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

SFP (8 x 40G + 96 x 10G)

When ports are in SFP mode, the mapping between physical ports and the logical ports/interfaces name is in the following table.

Physical Port number	L2/L3 port/interface name	interface support speed
1	qe-1/1/1	40Gb/s
2	qe-1/1/2	40Gb/s
3	qe-1/1/3	40Gb/s
4	qe-1/1/4	40Gb/s
5	4 x 10G	
	te-1/1/1	10Gb/s
	te-1/1/2	10Gb/s
	te-1/1/3	10Gb/s
	te-1/1/4	10Gb/s
6	4 x 10G	
	te-1/1/5	10Gb/s
	te-1/1/6	10Gb/s
	te-1/1/7	10Gb/s
	te-1/1/8	10Gb/s
7	4 x 10G	
	te-1/1/9	10Gb/s
	te-1/1/10	10Gb/s
	te-1/1/11	10Gb/s
	te-1/1/12	10Gb/s
8	4 x 10G	
	te-1/1/13	10Gb/s
	te-1/1/14	10Gb/s
	te-1/1/15	10Gb/s
	te-1/1/16	10Gb/s
9	4 x 10G	

	te-1/1/17	10Gb/s
	te-1/1/18	10Gb/s
	te-1/1/19	10Gb/s
	te-1/1/20	10Gb/s
10	4 x 10G	
	te-1/1/21	10Gb/s
	te-1/1/22	10Gb/s
	te-1/1/23	10Gb/s
	te-1/1/24	10Gb/s
11	4 x 10G	
	te-1/1/25	10Gb/s
	te-1/1/26	10Gb/s
	te-1/1/27	10Gb/s
	te-1/1/28	10Gb/s
12	4 x 10G	
	te-1/1/29	10Gb/s
	te-1/1/30	10Gb/s
	te-1/1/31	10Gb/s
	te-1/1/32	10Gb/s
13	4 x 10G	
	te-1/1/33	10Gb/s
	te-1/1/34	10Gb/s
	te-1/1/35	10Gb/s
	te-1/1/36	10Gb/s
14	4 x 10G	
	te-1/1/37	10Gb/s
	te-1/1/38	10Gb/s
	te-1/1/39	10Gb/s
	te-1/1/40	10Gb/s
15	4 x 10G	

	te-1/1/41	10Gb/s
	te-1/1/42	10Gb/s
	te-1/1/43	10Gb/s
	te-1/1/44	10Gb/s
16	4 x 10G	
	te-1/1/45	10Gb/s
	te-1/1/46	10Gb/s
	te-1/1/47	10Gb/s
	te-1/1/48	10Gb/s
17	qe-1/1/17	40Gb/s
18	qe-1/1/18	40Gb/s
19	qe-1/1/19	40Gb/s
20	qe-1/1/20	40Gb/s
21	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
22	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
23	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
24	4 x 10G	
	te-1/1/61	10Gb/s

	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
25	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
26	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s
27	4 x 10G	
	te-1/1/73	10Gb/s
	te-1/1/74	10Gb/s
	te-1/1/75	10Gb/s
	te-1/1/76	10Gb/s
28	4 x 10G	
	te-1/1/77	10Gb/s
	te-1/1/78	10Gb/s
	te-1/1/79	10Gb/s
	te-1/1/80	10Gb/s
29	4 x 10G	
	te-1/1/81	10Gb/s
	te-1/1/82	10Gb/s
	te-1/1/83	10Gb/s
	te-1/1/84	10Gb/s
30	4 x 10G	
	te-1/1/85	10Gb/s

	te-1/1/86	10Gb/s
	te-1/1/87	10Gb/s
	te-1/1/88	10Gb/s
31	4 x 10G	
	te-1/1/89	10Gb/s
	te-1/1/90	10Gb/s
	te-1/1/91	10Gb/s
	te-1/1/92	10Gb/s
32	4 x 10G	
	te-1/1/93	10Gb/s
	te-1/1/94	10Gb/s
	te-1/1/95	10Gb/s
	te-1/1/96	10Gb/s

40G Changes to 4*10G in L2/L3 mode on P-3922

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+4*40G)

When ports are in QSFP mode, the mapping between physical port, interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s

4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s

34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s

SFP (64*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s

7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s

37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s

	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s

40G Changes to 4*10G in L2/L3 mode on P-3920

In L2/L3 mode configuration

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+4*40G)

When ports are in QSFP mode, the mapping between physical port, interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s

15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s

45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s

SFP (64*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s

18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s

48	te-1/1/48	10Gb/s and 1Gb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s

40G Changes to 4*10G in L2/L3 mode on accton_as5712_54x/HP5712

In L2/L3 mode configuration

You can set the port mode by issuing:

```
admin@XorPlus# set interface ge-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+6*40G)

When ports are in QSFP mode, the mapping between physical port, interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s

24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s
53	qe-1/1/53	40Gb/s

54	qe-1/1/54	40Gb/s
----	-----------	--------

SFP (72*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s

26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	
	te-1/1/53	10Gb/s

	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
53	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
54	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s

40G Changes to 4*10G in L2/L3 mode on P-3930

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+4*40G)

When ports are in QSFP mode, the mapping between physical port, interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s , 1Gb/s and 100Mb/s
2	te-1/1/2	10Gb/s , 1Gb/s and 100Mb/s
3	te-1/1/3	10Gb/s , 1Gb/s and 100Mb/s
4	te-1/1/4	10Gb/s , 1Gb/s and 100Mb/s
5	te-1/1/5	10Gb/s , 1Gb/s and 100Mb/s
6	te-1/1/6	10Gb/s , 1Gb/s and 100Mb/s
7	te-1/1/7	10Gb/s , 1Gb/s and 100Mb/s
8	te-1/1/8	10Gb/s , 1Gb/s and 100Mb/s
9	te-1/1/9	10Gb/s , 1Gb/s and 100Mb/s
10	te-1/1/10	10Gb/s , 1Gb/s and 100Mb/s
11	te-1/1/11	10Gb/s , 1Gb/s and 100Mb/s
12	te-1/1/12	10Gb/s , 1Gb/s and 100Mb/s
13	te-1/1/13	10Gb/s , 1Gb/s and 100Mb/s
14	te-1/1/14	10Gb/s , 1Gb/s and 100Mb/s
15	te-1/1/15	10Gb/s , 1Gb/s and 100Mb/s
16	te-1/1/16	10Gb/s , 1Gb/s and 100Mb/s
17	te-1/1/17	10Gb/s , 1Gb/s and 100Mb/s
18	te-1/1/18	10Gb/s , 1Gb/s and 100Mb/s
19	te-1/1/19	10Gb/s , 1Gb/s and 100Mb/s
20	te-1/1/20	10Gb/s , 1Gb/s and 100Mb/s
21	te-1/1/21	10Gb/s , 1Gb/s and 100Mb/s
22	te-1/1/22	10Gb/s , 1Gb/s and 100Mb/s

23	te-1/1/23	10Gb/s , 1Gb/s and 100Mb/s
24	te-1/1/24	10Gb/s , 1Gb/s and 100Mb/s
25	te-1/1/25	10Gb/s , 1Gb/s and 100Mb/s
26	te-1/1/26	10Gb/s , 1Gb/s and 100Mb/s
27	te-1/1/27	10Gb/s , 1Gb/s and 100Mb/s
28	te-1/1/28	10Gb/s , 1Gb/s and 100Mb/s
29	te-1/1/29	10Gb/s , 1Gb/s and 100Mb/s
30	te-1/1/30	10Gb/s , 1Gb/s and 100Mb/s
31	te-1/1/31	10Gb/s , 1Gb/s and 100Mb/s
32	te-1/1/32	10Gb/s , 1Gb/s and 100Mb/s
33	te-1/1/33	10Gb/s , 1Gb/s and 100Mb/s
34	te-1/1/34	10Gb/s , 1Gb/s and 100Mb/s
35	te-1/1/35	10Gb/s , 1Gb/s and 100Mb/s
36	te-1/1/36	10Gb/s , 1Gb/s and 100Mb/s
37	te-1/1/37	10Gb/s , 1Gb/s and 100Mb/s
38	te-1/1/38	10Gb/s , 1Gb/s and 100Mb/s
39	te-1/1/39	10Gb/s , 1Gb/s and 100Mb/s
40	te-1/1/40	10Gb/s , 1Gb/s and 100Mb/s
41	te-1/1/41	10Gb/s , 1Gb/s and 100Mb/s
42	te-1/1/42	10Gb/s , 1Gb/s and 100Mb/s
43	te-1/1/43	10Gb/s , 1Gb/s and 100Mb/s
44	te-1/1/44	10Gb/s , 1Gb/s and 100Mb/s
45	te-1/1/45	10Gb/s , 1Gb/s and 100Mb/s
46	te-1/1/46	10Gb/s , 1Gb/s and 100Mb/s
47	te-1/1/47	10Gb/s , 1Gb/s and 100Mb/s
48	te-1/1/48	10Gb/s , 1Gb/s and 100Mb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s

SFP (64*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s , 1Gb/s and 100Mb/s
2	te-1/1/2	10Gb/s , 1Gb/s and 100Mb/s
3	te-1/1/3	10Gb/s , 1Gb/s and 100Mb/s
4	te-1/1/4	10Gb/s , 1Gb/s and 100Mb/s
5	te-1/1/5	10Gb/s , 1Gb/s and 100Mb/s
6	te-1/1/6	10Gb/s , 1Gb/s and 100Mb/s
7	te-1/1/7	10Gb/s , 1Gb/s and 100Mb/s
8	te-1/1/8	10Gb/s , 1Gb/s and 100Mb/s
9	te-1/1/9	10Gb/s , 1Gb/s and 100Mb/s
10	te-1/1/10	10Gb/s , 1Gb/s and 100Mb/s
11	te-1/1/11	10Gb/s , 1Gb/s and 100Mb/s
12	te-1/1/12	10Gb/s , 1Gb/s and 100Mb/s
13	te-1/1/13	10Gb/s , 1Gb/s and 100Mb/s
14	te-1/1/14	10Gb/s , 1Gb/s and 100Mb/s
15	te-1/1/15	10Gb/s , 1Gb/s and 100Mb/s
16	te-1/1/16	10Gb/s , 1Gb/s and 100Mb/s
17	te-1/1/17	10Gb/s , 1Gb/s and 100Mb/s
18	te-1/1/18	10Gb/s , 1Gb/s and 100Mb/s
19	te-1/1/19	10Gb/s , 1Gb/s and 100Mb/s
20	te-1/1/20	10Gb/s , 1Gb/s and 100Mb/s
21	te-1/1/21	10Gb/s , 1Gb/s and 100Mb/s
22	te-1/1/22	10Gb/s , 1Gb/s and 100Mb/s
23	te-1/1/23	10Gb/s , 1Gb/s and 100Mb/s
24	te-1/1/24	10Gb/s , 1Gb/s and 100Mb/s
25	te-1/1/25	10Gb/s , 1Gb/s and 100Mb/s
26	te-1/1/26	10Gb/s , 1Gb/s and 100Mb/s

27	te-1/1/27	10Gb/s , 1Gb/s and 100Mb/s
28	te-1/1/28	10Gb/s , 1Gb/s and 100Mb/s
29	te-1/1/29	10Gb/s , 1Gb/s and 100Mb/s
30	te-1/1/30	10Gb/s , 1Gb/s and 100Mb/s
31	te-1/1/31	10Gb/s , 1Gb/s and 100Mb/s
32	te-1/1/32	10Gb/s , 1Gb/s and 100Mb/s
33	te-1/1/33	10Gb/s , 1Gb/s and 100Mb/s
34	te-1/1/34	10Gb/s , 1Gb/s and 100Mb/s
35	te-1/1/35	10Gb/s , 1Gb/s and 100Mb/s
36	te-1/1/36	10Gb/s , 1Gb/s and 100Mb/s
37	te-1/1/37	10Gb/s , 1Gb/s and 100Mb/s
38	te-1/1/38	10Gb/s , 1Gb/s and 100Mb/s
39	te-1/1/39	10Gb/s , 1Gb/s and 100Mb/s
40	te-1/1/40	10Gb/s , 1Gb/s and 100Mb/s
41	te-1/1/41	10Gb/s , 1Gb/s and 100Mb/s
42	te-1/1/42	10Gb/s , 1Gb/s and 100Mb/s
43	te-1/1/43	10Gb/s , 1Gb/s and 100Mb/s
44	te-1/1/44	10Gb/s , 1Gb/s and 100Mb/s
45	te-1/1/45	10Gb/s , 1Gb/s and 100Mb/s
46	te-1/1/46	10Gb/s , 1Gb/s and 100Mb/s
47	te-1/1/47	10Gb/s , 1Gb/s and 100Mb/s
48	te-1/1/48	10Gb/s , 1Gb/s and 100Mb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s

	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s

40G Changes to 4*10G in L2/L3 mode on Niagara2632XL

In L2/L3 mode configuration

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/SFP-64/QSFP.
admin@XorPlus# commit
```

After setting ports to different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP (32 x 40G)

When ports are in QSFP mode, the mapping between physical ports, interface name, and interface support speed are in the following table.

Physical Port number	Interface name	Interface support speed
1	qe-1/1/1	40Gb/s
2	qe-1/1/2	40Gb/s
3	qe-1/1/3	40Gb/s
4	qe-1/1/4	40Gb/s

5	qe-1/1/5	40Gb/s
6	qe-1/1/6	40Gb/s
7	qe-1/1/7	40Gb/s
8	qe-1/1/8	40Gb/s
9	qe-1/1/9	40Gb/s
10	qe-1/1/10	40Gb/s
11	qe-1/1/11	40Gb/s
12	qe-1/1/12	40Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	qe-1/1/17	40Gb/s
18	qe-1/1/18	40Gb/s
19	qe-1/1/19	40Gb/s
20	qe-1/1/20	40Gb/s
21	qe-1/1/21	40Gb/s
22	qe-1/1/22	40Gb/s
23	qe-1/1/23	40Gb/s
24	qe-1/1/24	40Gb/s
25	qe-1/1/25	40Gb/s
26	qe-1/1/26	40Gb/s
27	qe-1/1/27	40Gb/s
28	qe-1/1/28	40Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

SFP-64 (16 x 40G + 64 x 10G)

When ports are in SFP-64 mode, the mapping between physical ports, interface name, and interface support speed is in the following table.

Physical Port number	Interface name	Interface support speed
1	4 x 10G	
	te-1/1/1	10Gb/s
	te-1/1/2	10Gb/s
	te-1/1/3	10Gb/s
	te-1/1/4	10Gb/s
2	4 x 10G	
	te-1/1/5	10Gb/s
	te-1/1/6	10Gb/s
	te-1/1/7	10Gb/s
	te-1/1/8	10Gb/s
3	4 x 10G	
	te-1/1/9	10Gb/s
	te-1/1/10	10Gb/s
	te-1/1/11	10Gb/s
	te-1/1/12	10Gb/s
4	4 x 10G	
	te-1/1/13	10Gb/s
	te-1/1/14	10Gb/s
	te-1/1/15	10Gb/s
	te-1/1/16	10Gb/s
5	4 x 10G	
	te-1/1/17	10Gb/s
	te-1/1/18	10Gb/s
	te-1/1/19	10Gb/s
	te-1/1/20	10Gb/s

6	4 x 10G	
	te-1/1/21	10Gb/s
	te-1/1/22	10Gb/s
	te-1/1/23	10Gb/s
	te-1/1/24	10Gb/s
7	4 x 10G	
	te-1/1/25	10Gb/s
	te-1/1/26	10Gb/s
	te-1/1/27	10Gb/s
	te-1/1/28	10Gb/s
8	4 x 10G	
	te-1/1/29	10Gb/s
	te-1/1/30	10Gb/s
	te-1/1/31	10Gb/s
	te-1/1/32	10Gb/s
9	qe-1/1/9	40Gb/s
10	qe-1/1/10	40Gb/s
11	qe-1/1/11	40Gb/s
12	qe-1/1/12	40Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	4 x 10G	
	te-1/1/33	10Gb/s
	te-1/1/34	10Gb/s
	te-1/1/35	10Gb/s
	te-1/1/36	10Gb/s
18	4 x 10G	
	te-1/1/37	10Gb/s

	te-1/1/38	10Gb/s
	te-1/1/39	10Gb/s
	te-1/1/40	10Gb/s
19	4 x 10G	
	te-1/1/41	10Gb/s
	te-1/1/42	10Gb/s
	te-1/1/43	10Gb/s
	te-1/1/44	10Gb/s
20	4 x 10G	
	te-1/1/45	10Gb/s
	te-1/1/46	10Gb/s
	te-1/1/47	10Gb/s
	te-1/1/48	10Gb/s
21	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
22	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
23	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
24	4 x 10G	
	te-1/1/61	10Gb/s

	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
25	qe-1/1/25	40Gb/s
26	qe-1/1/26	40Gb/s
27	qe-1/1/27	40Gb/s
28	qe-1/1/28	40Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

SFP (8 x 40G + 96 x 10G)

When ports are in SFP mode, the mapping between physical port number, interfaces name, and interface support speed is in the following table.

Physical Port number	Interface name	Interface support speed
1	4 x 10G	
	te-1/1/1	10Gb/s
	te-1/1/2	10Gb/s
	te-1/1/3	10Gb/s
	te-1/1/4	10Gb/s
2	4 x 10G	
	te-1/1/5	10Gb/s
	te-1/1/6	10Gb/s
	te-1/1/7	10Gb/s
	te-1/1/8	10Gb/s
3	4 x 10G	
	te-1/1/9	10Gb/s
	te-1/1/10	10Gb/s
	te-1/1/11	10Gb/s

	te-1/1/12	10Gb/s
4	4 x 10G	
	te-1/1/13	10Gb/s
	te-1/1/14	10Gb/s
	te-1/1/15	10Gb/s
	te-1/1/16	10Gb/s
5	4 x 10G	
	te-1/1/17	10Gb/s
	te-1/1/18	10Gb/s
	te-1/1/19	10Gb/s
	te-1/1/20	10Gb/s
6	4 x 10G	
	te-1/1/21	10Gb/s
	te-1/1/22	10Gb/s
	te-1/1/23	10Gb/s
	te-1/1/24	10Gb/s
7	4 x 10G	
	te-1/1/25	10Gb/s
	te-1/1/26	10Gb/s
	te-1/1/27	10Gb/s
	te-1/1/28	10Gb/s
8	4 x 10G	
	te-1/1/29	10Gb/s
	te-1/1/30	10Gb/s
	te-1/1/31	10Gb/s
	te-1/1/32	10Gb/s
9	4 x 10G	
	te-1/1/33	10Gb/s
	te-1/1/34	10Gb/s
	te-1/1/35	10Gb/s

	te-1/1/36	10Gb/s
10	4 x 10G	
	te-1/1/37	10Gb/s
	te-1/1/38	10Gb/s
	te-1/1/39	10Gb/s
	te-1/1/40	10Gb/s
11	4 x 10G	
	te-1/1/41	10Gb/s
	te-1/1/42	10Gb/s
	te-1/1/43	10Gb/s
	te-1/1/44	10Gb/s
12	4 x 10G	
	te-1/1/45	10Gb/s
	te-1/1/46	10Gb/s
	te-1/1/47	10Gb/s
	te-1/1/48	10Gb/s
13	qe-1/1/13	40Gb/s
14	qe-1/1/14	40Gb/s
15	qe-1/1/15	40Gb/s
16	qe-1/1/16	40Gb/s
17	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
18	4 x 10G	
	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s

19	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
20	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
21	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
22	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s
23	4 x 10G	
	te-1/1/73	10Gb/s
	te-1/1/74	10Gb/s
	te-1/1/75	10Gb/s
	te-1/1/76	10Gb/s
24	4 x 10G	
	te-1/1/77	10Gb/s
	te-1/1/78	10Gb/s
	te-1/1/79	10Gb/s
	te-1/1/80	10Gb/s

25	4 x 10G	
	te-1/1/81	10Gb/s
	te-1/1/82	10Gb/s
	te-1/1/83	10Gb/s
	te-1/1/84	10Gb/s
26	4 x 10G	
	te-1/1/85	10Gb/s
	te-1/1/86	10Gb/s
	te-1/1/87	10Gb/s
	te-1/1/88	10Gb/s
27	4 x 10G	
	te-1/1/89	10Gb/s
	te-1/1/90	10Gb/s
	te-1/1/91	10Gb/s
	te-1/1/92	10Gb/s
28	4 x 10G	
	te-1/1/93	10Gb/s
	te-1/1/94	10Gb/s
	te-1/1/95	10Gb/s
	te-1/1/96	10Gb/s
29	qe-1/1/29	40Gb/s
30	qe-1/1/30	40Gb/s
31	qe-1/1/31	40Gb/s
32	qe-1/1/32	40Gb/s

In L2/L3 mode configuration

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+6*40G)

When ports are in QSFP mode, the mapping between physical port, interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s

24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s
53	qe-1/1/53	40Gb/s

54	qe-1/1/54	40Gb/s
----	-----------	--------

SFP (72*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names, and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s

26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	
	te-1/1/53	10Gb/s

	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
53	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
54	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s

40G Changes to 4*10G in L2/L3 mode on as5812_54t

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state to take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+8*40G)

When ports are in QSFP mode, the mapping between physical port, interface names and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s , 1Gb/s and 100Mb/s
2	te-1/1/2	10Gb/s , 1Gb/s and 100Mb/s
3	te-1/1/3	10Gb/s , 1Gb/s and 100Mb/s
4	te-1/1/4	10Gb/s , 1Gb/s and 100Mb/s
5	te-1/1/5	10Gb/s , 1Gb/s and 100Mb/s
6	te-1/1/6	10Gb/s , 1Gb/s and 100Mb/s
7	te-1/1/7	10Gb/s , 1Gb/s and 100Mb/s
8	te-1/1/8	10Gb/s , 1Gb/s and 100Mb/s
9	te-1/1/9	10Gb/s , 1Gb/s and 100Mb/s
10	te-1/1/10	10Gb/s , 1Gb/s and 100Mb/s
11	te-1/1/11	10Gb/s , 1Gb/s and 100Mb/s
12	te-1/1/12	10Gb/s , 1Gb/s and 100Mb/s
13	te-1/1/13	10Gb/s , 1Gb/s and 100Mb/s
14	te-1/1/14	10Gb/s , 1Gb/s and 100Mb/s
15	te-1/1/15	10Gb/s , 1Gb/s and 100Mb/s
16	te-1/1/16	10Gb/s , 1Gb/s and 100Mb/s
17	te-1/1/17	10Gb/s , 1Gb/s and 100Mb/s
18	te-1/1/18	10Gb/s , 1Gb/s and 100Mb/s
19	te-1/1/19	10Gb/s , 1Gb/s and 100Mb/s
20	te-1/1/20	10Gb/s , 1Gb/s and 100Mb/s
21	te-1/1/21	10Gb/s , 1Gb/s and 100Mb/s
22	te-1/1/22	10Gb/s , 1Gb/s and 100Mb/s

23	te-1/1/23	10Gb/s , 1Gb/s and 100Mb/s
24	te-1/1/24	10Gb/s , 1Gb/s and 100Mb/s
25	te-1/1/25	10Gb/s , 1Gb/s and 100Mb/s
26	te-1/1/26	10Gb/s , 1Gb/s and 100Mb/s
27	te-1/1/27	10Gb/s , 1Gb/s and 100Mb/s
28	te-1/1/28	10Gb/s , 1Gb/s and 100Mb/s
29	te-1/1/29	10Gb/s , 1Gb/s and 100Mb/s
30	te-1/1/30	10Gb/s , 1Gb/s and 100Mb/s
31	te-1/1/31	10Gb/s , 1Gb/s and 100Mb/s
32	te-1/1/32	10Gb/s , 1Gb/s and 100Mb/s
33	te-1/1/33	10Gb/s , 1Gb/s and 100Mb/s
34	te-1/1/34	10Gb/s , 1Gb/s and 100Mb/s
35	te-1/1/35	10Gb/s , 1Gb/s and 100Mb/s
36	te-1/1/36	10Gb/s , 1Gb/s and 100Mb/s
37	te-1/1/37	10Gb/s , 1Gb/s and 100Mb/s
38	te-1/1/38	10Gb/s , 1Gb/s and 100Mb/s
39	te-1/1/39	10Gb/s , 1Gb/s and 100Mb/s
40	te-1/1/40	10Gb/s , 1Gb/s and 100Mb/s
41	te-1/1/41	10Gb/s , 1Gb/s and 100Mb/s
42	te-1/1/42	10Gb/s , 1Gb/s and 100Mb/s
43	te-1/1/43	10Gb/s , 1Gb/s and 100Mb/s
44	te-1/1/44	10Gb/s , 1Gb/s and 100Mb/s
45	te-1/1/45	10Gb/s , 1Gb/s and 100Mb/s
46	te-1/1/46	10Gb/s , 1Gb/s and 100Mb/s
47	te-1/1/47	10Gb/s , 1Gb/s and 100Mb/s
48	te-1/1/48	10Gb/s , 1Gb/s and 100Mb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s

53	qe-1/1/53	40Gb/s
54	qe-1/1/54	40Gb/s

SFP (72*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s , 1Gb/s and 100Mb/s
2	te-1/1/2	10Gb/s , 1Gb/s and 100Mb/s
3	te-1/1/3	10Gb/s , 1Gb/s and 100Mb/s
4	te-1/1/4	10Gb/s , 1Gb/s and 100Mb/s
5	te-1/1/5	10Gb/s , 1Gb/s and 100Mb/s
6	te-1/1/6	10Gb/s , 1Gb/s and 100Mb/s
7	te-1/1/7	10Gb/s , 1Gb/s and 100Mb/s
8	te-1/1/8	10Gb/s , 1Gb/s and 100Mb/s
9	te-1/1/9	10Gb/s , 1Gb/s and 100Mb/s
10	te-1/1/10	10Gb/s , 1Gb/s and 100Mb/s
11	te-1/1/11	10Gb/s , 1Gb/s and 100Mb/s
12	te-1/1/12	10Gb/s , 1Gb/s and 100Mb/s
13	te-1/1/13	10Gb/s , 1Gb/s and 100Mb/s
14	te-1/1/14	10Gb/s , 1Gb/s and 100Mb/s
15	te-1/1/15	10Gb/s , 1Gb/s and 100Mb/s
16	te-1/1/16	10Gb/s , 1Gb/s and 100Mb/s
17	te-1/1/17	10Gb/s , 1Gb/s and 100Mb/s
18	te-1/1/18	10Gb/s , 1Gb/s and 100Mb/s
19	te-1/1/19	10Gb/s , 1Gb/s and 100Mb/s
20	te-1/1/20	10Gb/s , 1Gb/s and 100Mb/s
21	te-1/1/21	10Gb/s , 1Gb/s and 100Mb/s
22	te-1/1/22	10Gb/s , 1Gb/s and 100Mb/s
23	te-1/1/23	10Gb/s , 1Gb/s and 100Mb/s
24	te-1/1/24	10Gb/s , 1Gb/s and 100Mb/s

25	te-1/1/25	10Gb/s , 1Gb/s and 100Mb/s
26	te-1/1/26	10Gb/s , 1Gb/s and 100Mb/s
27	te-1/1/27	10Gb/s , 1Gb/s and 100Mb/s
28	te-1/1/28	10Gb/s , 1Gb/s and 100Mb/s
29	te-1/1/29	10Gb/s , 1Gb/s and 100Mb/s
30	te-1/1/30	10Gb/s , 1Gb/s and 100Mb/s
31	te-1/1/31	10Gb/s , 1Gb/s and 100Mb/s
32	te-1/1/32	10Gb/s , 1Gb/s and 100Mb/s
33	te-1/1/33	10Gb/s , 1Gb/s and 100Mb/s
34	te-1/1/34	10Gb/s , 1Gb/s and 100Mb/s
35	te-1/1/35	10Gb/s , 1Gb/s and 100Mb/s
36	te-1/1/36	10Gb/s , 1Gb/s and 100Mb/s
37	te-1/1/37	10Gb/s , 1Gb/s and 100Mb/s
38	te-1/1/38	10Gb/s , 1Gb/s and 100Mb/s
39	te-1/1/39	10Gb/s , 1Gb/s and 100Mb/s
40	te-1/1/40	10Gb/s , 1Gb/s and 100Mb/s
41	te-1/1/41	10Gb/s , 1Gb/s and 100Mb/s
42	te-1/1/42	10Gb/s , 1Gb/s and 100Mb/s
43	te-1/1/43	10Gb/s , 1Gb/s and 100Mb/s
44	te-1/1/44	10Gb/s , 1Gb/s and 100Mb/s
45	te-1/1/45	10Gb/s , 1Gb/s and 100Mb/s
46	te-1/1/46	10Gb/s , 1Gb/s and 100Mb/s
47	te-1/1/47	10Gb/s , 1Gb/s and 100Mb/s
48	te-1/1/48	10Gb/s , 1Gb/s and 100Mb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	

	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
53	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
54	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s

40G Changes to 4*10G in L2/L3 mode on arctica4806xp

In L2/L3 mode configure

You can set the port mode by issuing:

```
admin@XorPlus# set interface qe-interface-mode SFP/QSFP.
admin@XorPlus# commit
```

After setting ports to a different mode, it is mandatory to restart the L2/L3 service in order to make the new state to take effect.

```
admin@XorPlus# run request system reboot
```

QSFP(48*10G+6*40G)

When ports are in QSFP mode, the mapping between physical port, interface names and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s

23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s
25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	qe-1/1/49	40Gb/s
50	qe-1/1/50	40Gb/s
51	qe-1/1/51	40Gb/s
52	qe-1/1/52	40Gb/s

53	qe-1/1/53	40Gb/s
54	qe-1/1/54	40Gb/s

SFP (72*10G)

When ports are in SFP mode, the mapping between physical port, the logical ports/interface names and interfaces support speed is in the following table.

Physical Port number	interface name	interface support speed
1	te-1/1/1	10Gb/s and 1Gb/s
2	te-1/1/2	10Gb/s and 1Gb/s
3	te-1/1/3	10Gb/s and 1Gb/s
4	te-1/1/4	10Gb/s and 1Gb/s
5	te-1/1/5	10Gb/s and 1Gb/s
6	te-1/1/6	10Gb/s and 1Gb/s
7	te-1/1/7	10Gb/s and 1Gb/s
8	te-1/1/8	10Gb/s and 1Gb/s
9	te-1/1/9	10Gb/s and 1Gb/s
10	te-1/1/10	10Gb/s and 1Gb/s
11	te-1/1/11	10Gb/s and 1Gb/s
12	te-1/1/12	10Gb/s and 1Gb/s
13	te-1/1/13	10Gb/s and 1Gb/s
14	te-1/1/14	10Gb/s and 1Gb/s
15	te-1/1/15	10Gb/s and 1Gb/s
16	te-1/1/16	10Gb/s and 1Gb/s
17	te-1/1/17	10Gb/s and 1Gb/s
18	te-1/1/18	10Gb/s and 1Gb/s
19	te-1/1/19	10Gb/s and 1Gb/s
20	te-1/1/20	10Gb/s and 1Gb/s
21	te-1/1/21	10Gb/s and 1Gb/s
22	te-1/1/22	10Gb/s and 1Gb/s
23	te-1/1/23	10Gb/s and 1Gb/s
24	te-1/1/24	10Gb/s and 1Gb/s

25	te-1/1/25	10Gb/s and 1Gb/s
26	te-1/1/26	10Gb/s and 1Gb/s
27	te-1/1/27	10Gb/s and 1Gb/s
28	te-1/1/28	10Gb/s and 1Gb/s
29	te-1/1/29	10Gb/s and 1Gb/s
30	te-1/1/30	10Gb/s and 1Gb/s
31	te-1/1/31	10Gb/s and 1Gb/s
32	te-1/1/32	10Gb/s and 1Gb/s
33	te-1/1/33	10Gb/s and 1Gb/s
34	te-1/1/34	10Gb/s and 1Gb/s
35	te-1/1/35	10Gb/s and 1Gb/s
36	te-1/1/36	10Gb/s and 1Gb/s
37	te-1/1/37	10Gb/s and 1Gb/s
38	te-1/1/38	10Gb/s and 1Gb/s
39	te-1/1/39	10Gb/s and 1Gb/s
40	te-1/1/40	10Gb/s and 1Gb/s
41	te-1/1/41	10Gb/s and 1Gb/s
42	te-1/1/42	10Gb/s and 1Gb/s
43	te-1/1/43	10Gb/s and 1Gb/s
44	te-1/1/44	10Gb/s and 1Gb/s
45	te-1/1/45	10Gb/s and 1Gb/s
46	te-1/1/46	10Gb/s and 1Gb/s
47	te-1/1/47	10Gb/s and 1Gb/s
48	te-1/1/48	10Gb/s and 1Gb/s
49	4 x 10G	
	te-1/1/49	10Gb/s
	te-1/1/50	10Gb/s
	te-1/1/51	10Gb/s
	te-1/1/52	10Gb/s
50	4 x 10G	

	te-1/1/53	10Gb/s
	te-1/1/54	10Gb/s
	te-1/1/55	10Gb/s
	te-1/1/56	10Gb/s
51	4 x 10G	
	te-1/1/57	10Gb/s
	te-1/1/58	10Gb/s
	te-1/1/59	10Gb/s
	te-1/1/60	10Gb/s
52	4 x 10G	
	te-1/1/61	10Gb/s
	te-1/1/62	10Gb/s
	te-1/1/63	10Gb/s
	te-1/1/64	10Gb/s
53	4 x 10G	
	te-1/1/65	10Gb/s
	te-1/1/66	10Gb/s
	te-1/1/67	10Gb/s
	te-1/1/68	10Gb/s
54	4 x 10G	
	te-1/1/69	10Gb/s
	te-1/1/70	10Gb/s
	te-1/1/71	10Gb/s
	te-1/1/72	10Gb/s

LLDP Configuration (Link Layer Discovery Protocol)

LLDP is a standard link-layer discovery protocol which can broadcast its capability, IP address, ID, and interface name as TLVs (Type/Length/Value) in LLDP PDUs (Link Layer Discovery Protocol Data Units).

An LLDP PDU includes 4 basic TLVs and several optional TLVs.

Basic TLVs include the Chassis ID, Port ID, TTL and End TLVs.

In L2/L3, you can select the following optional TLVs:

TLV Name	Description
mac-phy-cfg	MAC address of the system
management-address	Management IP address of the system
port-description	The port description of system
port-vlan	The VLAN ID of the port
system-capabilities	System capability (e.g. switching, routing)
system-description	System description
system-name	System name

Configuring the LLDP Mode

LLDP supports 4 modes: TxRx, Tx_only, Rx_only, and Disabled. In TxRx mode, the system transmits *and* receives LLDPDUs. In Tx_only, the system only transmits LLDPDUs. In Rx_only, the system only receives LLDPDUs. In Disabled, the system will not transmit or receive any LLDPDUs.

You can configure the system as shown below:

```
admin@XorPlus# set protocols lldp enable true
admin@XorPlus# set protocols lldp interface te-1/1/1 working-mode tx_rx
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Selecting Optional TLVs

```
admin@XorPlus# set protocols lldp tlv-select mac-phy-cfg true
admin@XorPlus# set protocols lldp tlv-select management-address true
admin@XorPlus# set protocols lldp tlv-select port-description true
admin@XorPlus# set protocols lldp tlv-select system-capabilities true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Displaying LLDP Information

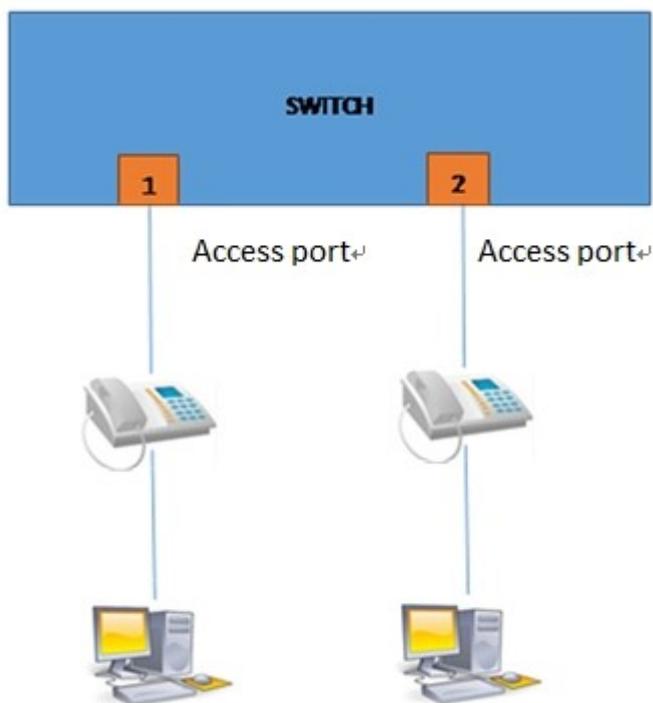
```
admin@XorPlus# show protocols lldp
Waiting for building configuration.
enable: true
tlv-select {
}
```

Configuring Other Parameters

You can configure other parameters, such as advertisement-interval, hold-time-multiplier, reinit-delay, and transmit-delay, in a similar manner.

Voice Vlan with LLDP Configuration Guide

Voice vlan with LLDP means pica8 switch will learn OUI address from the source mac address of the LLDP packet if LLDP is enabled between ip phone and pica8 switch.



Basic Configuration:

Step1.Enable LLDP

```
set protocols lldp enable true
```

Step2. Configuration voice vlan

- Configure the auto mode of voice vlan on the access port:

```
set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode auto
```

- Configure the manual mode of voice vlan on the access port:

```

set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode manual

```

Step3. Configuration Output port

```

set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan members 10

```

Configurations example

1. Access port: Port Ge-1/1/45 and Port Ge-1/1/46 are configured voice-vlan 100 in the manual mode.
2. Access port: Ge-1/1/47 and Ge-1/1/48 are configured voice-vlan 200 in the auto mode(default).
3. Output port : Ge-1/1/1 are configured vlan member 100 and 200 in the trunk mode.

```

admin@XorPlus# show | display set
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/45 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/45 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/46 family ethernet-switching native-vlan-id 1001
set interface gigabit-ethernet ge-1/1/46 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/46 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/46 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching native-vlan-id 1002
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/47 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching native-vlan-id 1003
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/48 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 100
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 200
set protocols lldp enable true
set vlans vlan-id 100
set vlans vlan-id 200
set vlans vlan-id 1000-1003

```

- If the ports configured voice-vlan received lldp message from the phone, show oui, we'll get lldp OUI addresses.

```

admin@XorPlus# run show vlans voice-vlan oui
Oui_Address      Mask      Description
0:1:e3:0:0:0     ff:ff:ff:0:0:0   Siemens phone
0:3:6b:0:0:0     ff:ff:ff:0:0:0   Cisco phone
0:4:d:0:0:0      ff:ff:ff:0:0:0   Avaya phone
0:60:b9:0:0:0    ff:ff:ff:0:0:0   Philips/NEC phone
0:d0:1e:0:0:0    ff:ff:ff:0:0:0   Pingtel phone
0:e0:75:0:0:0    ff:ff:ff:0:0:0   Polycom phone
0:e0:bb:0:0:0    ff:ff:ff:0:0:0   3com phone
Learned_Oui_Address  Mask
22:22:22:22:11:11    ff:ff:ff:ff:ff:ff
22:22:22:22:22:22    ff:ff:ff:ff:ff:ff
22:22:22:22:33:33    ff:ff:ff:ff:ff:ff
22:22:22:22:44:44    ff:ff:ff:ff:ff:ff

```

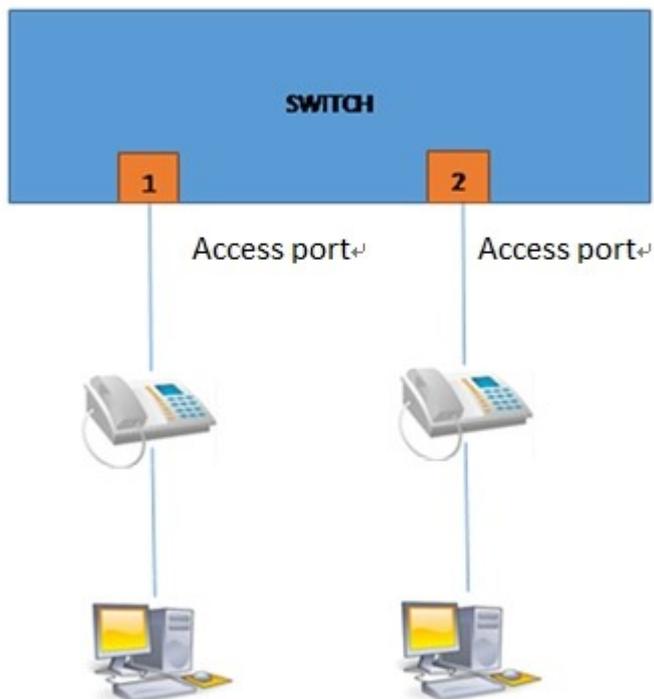
- Show voice-vlan, in manual mode, the lldp source mac address learned won't be shown in the table, while in auto mode, the lldp source mac address learned will be shown in the table. If an ingress port configure manual mode, as long as source mac address of phone can match any one of all static OUI and learned OUI, packets from phone can be tagged in voice vlan id. But if in auto mode, as long as source mac address of phone can only match OUI learned by the port, packets from phone can be tagged in voice vlan id.

```
admin@XorPlus# run show vlans voice-vlan vlan-id 100
Voice Vlan ID:100
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port          Mode      Tagged   Mac_Address      Status
-----
ge-1/1/45     manual   false
ge-1/1/46     manual   false
admin@XorPlus# run show vlans voice-vlan vlan-id 200
Voice Vlan ID:200
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port          Mode      Tagged   Mac_Address      Status
-----
ge-1/1/47     auto     false    22:22:22:22:33:33 Working
ge-1/1/48     auto     false    22:22:22:22:44:44 Working
```

Voice Vlan with LLDP Compliance CDP Configuration Guide

Voice vlan with LLDP Compliance CDP means the LLDP feature can get information from CDP packets (Cisco Discovery Protocol packets).

This feature is customized for Cisco ip phone.



Basic Configuration:

Step1.Enable LLDP Compliance CDP

- Enable LLDP Compliance CDP

```
set protocols lldp enable true
set protocols lldp compliance cdp true
set protocols lldp interface ge-1/1/1 compliance cdp true
set protocols lldp interface ge-1/1/2 compliance cdp true
```

- Check the configuration status

```
admin@XorPlus# show protocols lldp
enable: true
interface "ge-1/1/1" {
  compliance {
    cdp: true
  }
}
interface "ge-1/1/2" {
  compliance {
    cdp: true
  }
}
compliance {
  cdp: true
}

admin@XorPlus# run show lldp neighbor ge-1/1/1 detail
Local Port: ge-1/1/1
CDP info:
+Device ID: R1
+Address:
  Protocol type: NLPID format
  Protocol: IPv4
  Ip_address: 192.168.0.1
```

```

Protocol type: NLPID format
Protocol: IPv4
Ip_address: 192.168.0.2
+Port ID: GigabitEthernet0/5
+Capabilities:
  Router: YES
  Transparent Bridge: YES
  Source Route Bridge: YES
  Switch: YES
  Host: YES
  IGMP capable: YES
  Repeater: YES
+Version: Cisco IOS Software, 3700 Software (C3725-ADVENTERPRISEK9-M), Version 12.4(9)T1,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2006 by Cisco Systems, Inc.
Compiled Wed 30-Aug-06 18:48 by prod_rel_team
+Platform: cisco ip phone 7965
+Vtp_management_domain: Lab
+Native vlan: 1
+Full/Half Duplex: Full

The information above is just an example.

```

Step2. Configuration voice vlan

- Configure the auto mode of voice vlan on the access port:

```

set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode auto

```

- Configure the manual mode of voice vlan on the access port:

```

set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode manual

```

Step3. Configuration Output port

```

set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan members 10

```

Configurations example

1. Access port: Port Ge-1/1/45 and Port Ge-1/1/46 are configured voice-vlan 100 in the manual mode.
2. Access port: Ge-1/1/47 and Ge-1/1/48 are configured voice-vlan 200 in the auto mode(default).
3. Output port : Ge-1/1/1 are configured vlan member 100 and 200 in the trunk mode.

```

admin@XorPlus# show | display set
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/45 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/45 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/46 family ethernet-switching native-vlan-id 1001

```

```

set interface gigabit-ethernet ge-1/1/46 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/46 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/46 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching native-vlan-id 1002
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/47 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching native-vlan-id 1003
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/48 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 100
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 200
set protocols lldp enable true
set vlans vlan-id 100
set vlans vlan-id 200
set vlans vlan-id 1000-1003

```

- If the ports configured voice-vlan received CDP message from the IP phone, show oui, we'll get the learned OUI addresses that are source mac addresses in CDP packets.

```

admin@XorPlus# run show vlans voice-vlan oui
Oui_Address      Mask      Description
0:1:e3:0:0:0     ff:ff:ff:0:0:0   Siemens phone
0:3:6b:0:0:0     ff:ff:ff:0:0:0   Cisco phone
0:4:d:0:0:0      ff:ff:ff:0:0:0   Avaya phone
0:60:b9:0:0:0    ff:ff:ff:0:0:0   Philips/NEC phone
0:d0:1e:0:0:0    ff:ff:ff:0:0:0   Pingtel phone
0:e0:75:0:0:0    ff:ff:ff:0:0:0   Polycom phone
0:e0:bb:0:0:0    ff:ff:ff:0:0:0   3com phone
Learned_Oui_Address  Mask
22:22:22:22:11:11   ff:ff:ff:ff:ff:ff
22:22:22:22:22:22   ff:ff:ff:ff:ff:ff
22:22:22:22:33:33   ff:ff:ff:ff:ff:ff
22:22:22:22:44:44   ff:ff:ff:ff:ff:ff

```

- Show voice-vlan, in manual mode, the learned source mac address of CDP packet won't be shown in the table, while in auto mode, the learned source mac address of CDP packet will be shown in the table. If an ingress port configure manual mode, as long as source mac address of phone can match any one of all static OUI and learned OUI, packets from phone can be tagged in voice vlan id. But if in auto mode, as long as source mac address of phone can only match OUI learned by the port, packets from phone can be tagged in voice vlan id.

```

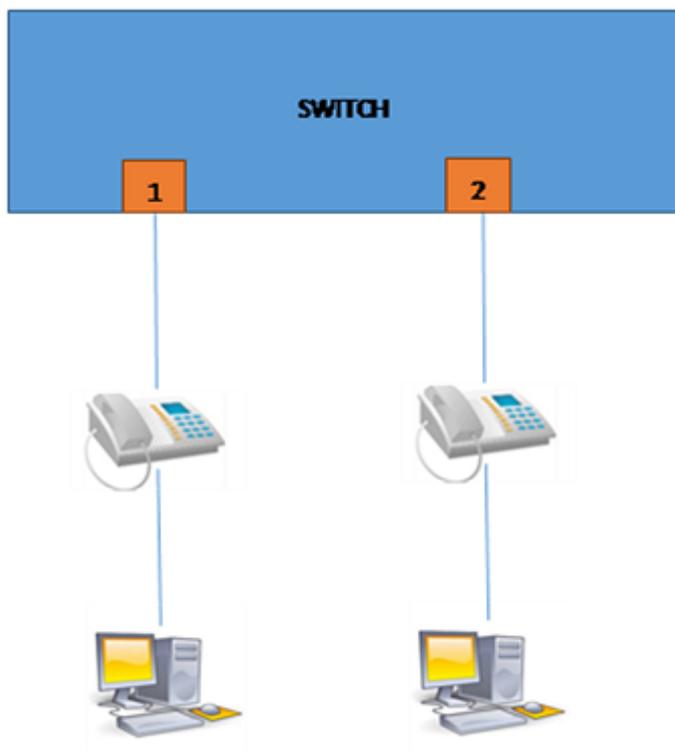
admin@XorPlus# run show vlans voice-vlan vlan-id 100
Voice Vlan ID:100
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port      Mode      Tagged   Mac_Address      Status
-----
ge-1/1/45  manual   false    Working
ge-1/1/46  manual   false    Working
admin@XorPlus# run show vlans voice-vlan vlan-id 200
Voice Vlan ID:200
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port      Mode      Tagged   Mac_Address      Status
-----
ge-1/1/47  auto     false    22:22:22:22:33:33  Working
ge-1/1/48  auto     false    22:22:22:22:44:44  Working

```

Voice Vlan Configuration Guide

Voice VLAN is a special VLAN designed to handle voice traffic. To improve the quality of voice traffic, traffics in voice VLAN is configured customized QoS (Quality of Service).

Switch could recognize voice traffic according to its source MAC address. If source MAC address of traffic matches OUI(Organizational Unique Identifier), it can be considered as voice traffic.



Default OUI:

- 1 0001-e300-0000 Siemens phone
- 2 0003-6b00-0000 Cisco phone
- 3 0004-0d00-0000 Avaya phone
- 4 0060-b900-0000 Philips/NEC phone
- 5 00d0-1e00-0000 Pingtel phone
- 6 00e0-7500-0000 Polycom phone
- 7 00e0-bb00-0000 3com phone

In our design, it is allowed to add more OUIs. However, 3 more OUI (10 in total) is the maximum number we support. Also, we support 4 VOICE VLANs at most. Each voice VLAN support 10 OUI at most.

There are two type of mode of voice VLAN configuration, that is, auto mode and manual mode.

Auto Mode:

Until learned mac address, auto mode will automatically add port into voice VLAN and install corresponding acl rules. Delete the port from the voice VLAN and uninstall acl rules after the voice-aging time since the mac address is aged.

Manual Mode:

Manual mode allows user manually add port into voice VLAN. When the learned OUI address is aged, the port will not be deleted from voice vlan, while it will be deleted in the auto mode.

- ✔ Voice VLAN could not be configured as pvid or vlan members on a port
Voice vlan NOT supported on LAG interface.
Voice vlan Must be configured in the port with trunk port-mode.

For the case, traffic from Computer and traffic from VOIP may all arrive at the same port. Since both of the traffics are all untagged packages, we use ACLs to separate these two traffics. Any untagged VOIP traffic will tagged the number of voice vlan. Other untagged traffic from computer is tagged the number its PVID. In this case, we can separate these two untagged traffic into two different vlans. And traffic from computer and from VOIP can work at the same time.

Basic Configuration

```
set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
```

Configuration Examples:

In Auto mode:

```
set vlans vlan-id 10
set vlans voice-vlan aging 600
set vlans voice-vlan mac-address 00:11:11:00:00:01 mask ff:ff:ff:00:00:00 description test
set vlans voice-vlan local-priority 6
set vlans voice-vlan dscp 63
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode auto
set interface gigabit-ethernet ge-1/1/1 voice-vlan tagged
```

In Manual mode:

```
set vlans vlan-id 10
set vlans voice-vlan aging 600
set vlans voice-vlan mac-address 00:11:11:00:00:01 mask ff:ff:ff:00:00:00 description test
set vlans voice-vlan local-priority 6
set vlans voice-vlan dscp 63
```

```
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode manual
set interface gigabit-ethernet ge-1/1/1 voice-vlan tagged
```

LLDP MED Configuration

In LLDP-MED protocol, it provides extra Organizationally Specific TLVs. The TLVs lists that Pica8 switch support are as follows :

M Mandatory for this TLV to be included in all outgoing LLDP-MED LLDPDUs,

C Conditionally required, mandatory for this TLV to be included in outgoing LLDP-MED LLDPDUs under the described conditions,

LLDP-MED TLV Subtype	TLV Name	LLDPDU Usage
1	LLDP-MED Capabilities	M
2	Network Policy	C1
3	Location Identification	C1
4	Extended Power-via-MDI	C2

Note 1: Transmission of these mandatory TLVs only applies if the associated TLV data contents have been administratively configured and apply on a given port.

Note 2: Extended Power-via-MDI TLV is mandatory for IEEE 802.3af compliant PSE Network Connectivity Devices. This TLV is not applicable for non-PoE capable ports, and should not be included in outgoing LLDP-MED LLDPDUs on such ports.

Basic Configuration:

Enable LLDP

```
set protocols lldp enable true
```

Configure lldp med tlvs

- basically these tlvs are selected by default:

```
set protocols lldp med-tlv-select inventory-management true/false
set protocols lldp med-tlv-select network-policy true/false
set protocols lldp med-tlv-select extended-power-via-MDI true/false
```

- Configure the lldp med fast start repeat count:
- The LLDP-MED fast start repeat count specifies the number of LLDP packets that will be sent during the LLDP-MED fast start period.

```
set protocols lldp med-fast-start-repeat-count 5
```

The LLDP MED function will only be active by receiving lldp med packets.

Show the med working status:

```
admin@XorPlus# run show lldp detail
LLDP: Enable
Advertisement interval: 30
Re-initialization Delay: 2
Transmit Delay: 2
Hold timer: 120
LLDP-MED fast start repeat count: 5
Selected TLVs:
  port_description
  system_name
  system_description
  system_capabilities
  management_address
  port_vlan_id
  mac_phy
Selected MED TLVs:
  inventory management
  network_policy
  extended_power_via_mdi
```

Show the med information of neighbor device:

```
admin@XorPlus# run show lldp neighbor ge-1/1/1 detail
Local Port: ge-1/1/1
LLDP info:
  Time To Live: 180
  Chassis Id: 192.1.1.1
  Port ID: 189C5DB7E4F4:P1
  Port Description: SW PORT
  System Name: SEP189C5DB7E4F4
  System Description: Cisco IP Phone 7965G,V14,
  System Capability: Bridge, Telephone
  Management Address:
  Default VLAN ID: 0
  Auto Negotiation: Supported, Enabled
  Physical media capabilities: FDX_S_Pause, FDX_B_Pause, 1000base_XFD, 1000base_T
  Media Attachment Unit type: 1000base_T_Full_Duplex
  +Med capabilities: Capabilities, Network Policy, Extended Power via MDI-PD, Inventory
  Med device type: Endpoint Class III
  +MED Network Policy
  Application Type: Voice
  Policy Flags: Known Policy
  Vlan ID: 4095
  L2 Priority: 5
  DSCP Value: 46
  +MED Network Policy
  Application Type: Voice Signaling
  Policy Flags: Known Policy
  Vlan ID: 4095
  L2 Priority: 0
  DSCP Value: 0
  +MED Extended Power via MDI
  Power Type: PD device
  Power Source: Unknown
```

```

Power Priority: Unknown
Power Value: 12.0 watts
+MED Hardware revision: 14
+MED Firmware revision: tnp65.9-3-1-CR17.bin
+MED Software revision:
+MED Serial number: FCH174499U2
+MED Manufacturer: Cisco Systems, Inc.
+MED Model name: CP-7965G
+MED Asset ID:

```

Show the med information of local device :

```

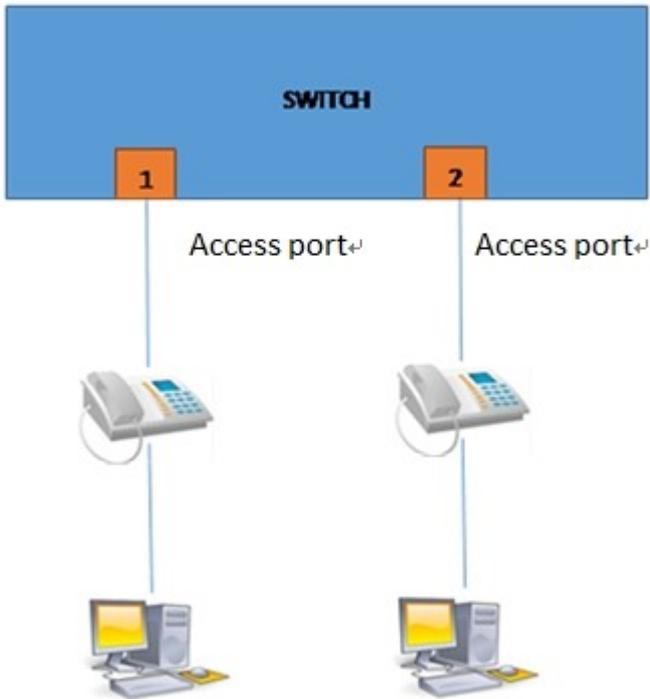
admin@XorPlus# run show lldp local_info ge-1/1/1 detail
LLDP Local configuration details
Chassis ID: 70:72:cf:fd:8f:21
System name: XorPlus
System description: PICA8 Inc., Model as4610_30p, PicOS 2.9.1
Interface LLDP State
-----
ge-1/1/1 Enable tx_rx

+Med capabilities: Capabilities, Network Policy, Extended Power via MDI-PD, Inventory
Med device type: Network Connectivity
+MED Network Policy
Application Type: Voice
Policy Flags: Unknown Policy
Vlan ID: 0
L2 Priority: 6
DSCP Value: 46
+MED Network Policy
Application Type: Voice Signaling
Policy Flags: Unknown Policy
Vlan ID: 0
L2 Priority: 0
DSCP Value: 0
+MED Extended Power via MDI
Power Type: PSE device
Power Source: Primary Power Source
Power Priority: Unknown
Power Value: 0.0 watts
+MED Hardware revision: N/A
+MED Firmware revision: 2.9.1/ac36038
+MED Software revision: 2.9.1/ac36038
+MED Serial number: AF10029779
+MED Manufacturer: Edgecore
+MED Model name: as4610_30p
+MED Asset ID: N/A

```

Voice Vlan with LLDP MED Configuration Guide

Voice vlan with LLDP MED means that pica8 switch will learn OUI address from the source mac address of the LLDP MED packet if LLDP is enabled between ip phone and pica8 switch.



Basic Configuration:

Step1.Enable LLDP

```
set protocols lldp enable true
```

Step2. Configuration voice vlan

- Configure the auto mode of voice vlan on the access port:

```
set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode auto
```

- Configure the manual mode of voice vlan on the access port:

```
set vlans vlan-id 10
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 voice-vlan vlan-id 10
set interface gigabit-ethernet ge-1/1/1 voice-vlan mode manual
```

Step3. Configuration Output port

```
set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan members 10
```

Configurations example

- 1.Access port: Port Ge-1/1/45 and Port Ge-1/1/46 are configured voice-vlan 100 in the manual mode.

2. Access port: Ge-1/1/47 and Ge-1/1/48 are configured voice-vlan 200 in the auto mode (default).

3. Output port : Ge-1/1/1 are configured vlan member 100 and 200 in the trunk mode.

```
admin@XorPlus# show | display set
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet ge-1/1/45 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/45 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/45 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/46 family ethernet-switching native-vlan-id 1001
set interface gigabit-ethernet ge-1/1/46 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/46 voice-vlan mode "manual"
set interface gigabit-ethernet ge-1/1/46 voice-vlan vlan-id 100
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching native-vlan-id 1002
set interface gigabit-ethernet ge-1/1/47 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/47 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching native-vlan-id 1003
set interface gigabit-ethernet ge-1/1/48 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/48 voice-vlan vlan-id 200
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 100
set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 200
set protocols lldp enable true
set vlans vlan-id 100
set vlans vlan-id 200
set vlans vlan-id 1000-1003
```

- If the ports configured voice-vlan received lldp med message from the phone, show oui, we'll get learned OUI addresses.

```
admin@XorPlus# run show vlans voice-vlan oui
Oui_Address      Mask      Description
0:1:e3:0:0:0     ff:ff:ff:0:0:0   Siemens phone
0:3:6b:0:0:0     ff:ff:ff:0:0:0   Cisco phone
0:4:d:0:0:0      ff:ff:ff:0:0:0   Avaya phone
0:60:b9:0:0:0    ff:ff:ff:0:0:0   Philips/NEC phone
0:d0:1e:0:0:0    ff:ff:ff:0:0:0   Pingtel phone
0:e0:75:0:0:0    ff:ff:ff:0:0:0   Polycom phone
0:e0:bb:0:0:0    ff:ff:ff:0:0:0   3com phone
Learned_Oui_Address  Mask
22:22:22:22:11:11   ff:ff:ff:ff:ff:ff
22:22:22:22:22:22   ff:ff:ff:ff:ff:ff
22:22:22:22:33:33   ff:ff:ff:ff:ff:ff
22:22:22:22:44:44   ff:ff:ff:ff:ff:ff
```

- Show voice-vlan, in manual mode, the lldp med source mac address learned won't be shown in the table, while in auto mode, the lldp med source mac address learned will be shown in the table. If an ingress port configure manual mode, as long as source mac address of phone can match any one of all static OUI and learned OUI, packets from phone can be tagged in voice vlan id. But if in auto mode, as long as source mac address of phone can only match OUI learned by the port, packets from phone can be tagged in voice vlan id.

```
admin@XorPlus# run show vlans voice-vlan vlan-id 100
Voice Vlan ID:100
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port      Mode      Tagged      Mac_Address      Status
-----
```

```

ge-1/1/45      manual    false
ge-1/1/46      manual    false
admin@XorPlus# run show vlans voice-vlan vlan-id 200
Voice Vlan ID:200
Voice Vlan local priority:6
Voice Vlan dscp:46
Voice Vlan aging time:1440 minutes
Current voice vlan enabled port mode:
Port          Mode      Tagged   Mac_Address      Status
-----
ge-1/1/47     auto     false   22:22:22:22:33:33 Working
ge-1/1/48     auto     false   22:22:22:22:44:44 Working

```

Static Link Aggregation (LAG) Configuration

- You can configure up to 48 LAGs in L2/L3, and each LAG can have up to 8 member ports.
- Both static and LACP LAGs can support the hashing of traffic using the Src/Dst MAC address, the Src/DstIP address, and Layer 4 port information.
- If all member ports of a LAN are link-down, the LAG will be link-down. The LAG will become link-up when at least one member port is link-up.
- The logical function and configuration of LAGs are same as those of a physical port.

Configuring Static LAGs

```

admin@XorPlus# set interface aggregate-ethernet ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 ether-options 802.3ad ae1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Displaying Static LAG Information

```

admin@XorPlus# run show interface aggregate-ethernet ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 53
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: c8:0a:a9:9e:14:9f, Hardware address: c8:0a:a9:9e:14:9f
Traffic statistics:
Input Packets.....176
Output Packets.....16
Input Octets.....12888
Output Octets.....1594
Aggregated link protocol: STATIC
Members Status Port Speed
-----
ge-1/1/1 Down Auto
ge-1/1/2 Down Auto
ge-1/1/3 Up Auto
ge-1/1/4 Up Auto

```

LAG Hashing Configuration

The IEEE 802.3ad link aggregation protocol groups multiple Ethernet interfaces and forms a single link layer interface known as LAG (link aggregation group). Traffic is balanced across the member links in the LAG to make use of all available bandwidth. The balancing is done by the LAG hashing algorithm. The LAG hashing algorithm determines the member link to be used for an incoming frame on the basis of certain values in the frame header.

- LAG Hashing Configuration and Example
- LAG Hash Mapping
- Resilient LAG Hashing Configuration and Example

LAG Hashing Configuration and Example

This example configures the LAG hash mapping mode as **ethernet-destination-only**.

Configuration:

```
set interface aggregate-ethernet ae1 hash-mapping mode ethernet-destination-only
```

Examples:

Configure one LAG with three ports.

```
set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/3 ether-options 802.3ad ae10
```

Configure the LAG hash mapping mode as **ethernet-destination-only**.

```
set interface aggregate-ethernet ae10 hash-mapping mode ethernet-destination-only
```

LAG Hash Mapping

User can configure the LAG hash mapping field according to the requirement.

Configure one LAG with three ports.

```
set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/3 ether-options 802.3ad ae10
```

Configure the LAG hash mapping mode as **ethernet-destination-only**.

```
set interface aggregate-ethernet ae10 hash-mapping mode ethernet-destination-only
```

Configure the LAG hash mapping field.

```
set interface aggregate-balancing hash-mapping field ethernet-destination-address disable
false
set interface aggregate-balancing hash-mapping field ethernet-source-address disable false
```

Resilient LAG Hashing Configuration and Example

Configure the LAG hashing mode as **advanced-resilient**. By default, the hash-mapping field is **disable false all**.

Without resilient mode, each traffic flow (the flow definition depends on the hash-mapping configuration) is load balanced on one port of the LAG. This distribution is done via a hashing algorithm. If a port on this LAG is added or removed (link up/down), the hash will change all flows to be re-distributed on the remaining ports. It is typically not important if the device on the other side of the LAG is a router or a switch but could be important if the LAG is sending traffic to equipment (like a cluster of servers) that handles traffic differently on each port (for example, distributing http sessions on multiple servers).

In resilient mode, a removed link does not trigger the redistribution of traffic on the remaining ports. Only the traffic from the removed port will be distributed to the remaining ports. An added link does not trigger redistribution of traffic on the remaining ports. The added link will share the traffic on the remaining ports.



Currently, all the switch models, except Tomahawk based switches, support the resilient LAG hashing.

Configuration:

```
set interface aggregate-ethernet ae10 hash-mapping mode advanced-resilient
```

Examples:

Configure one lag with three ports

```
set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae10
set interface gigabit-ethernet ge-1/1/3 ether-options 802.3ad ae10
```

Configure the lag hash mode advanced-resilient

```
set interface aggregate-ethernet ae10 hash-mapping mode advanced-resilient
```

Configure the hash-mapping field

```
set interface aggregate-balancing hash-mapping field ip-destination disable false
set interface aggregate-balancing hash-mapping field ip-source disable false
```

Link Aggregation Control Protocol (LACP) Configuration

LACP (802.3ad) provides the dynamic link aggregation function.

The LACPDU includes the LACP system priority, the system MAC, the port priority, and I.D. The port, included in the LACP LAG, will transmit the LACPDU to its neighbors.

The configuration of the LACP LAG is similar to that of the static LAG.

The **min-selected-port** label denotes that the LAG is up only when no fewer than the defined number of ports are up. Below, our defined number is 4.

Configuring LACP LAGs

```
admin@XorPlus# set interface aggregate-ethernet ael aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ael aggregated-ether-options
min-selected-port 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ael
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ael
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 ether-options 802.3ad ael
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 ether-options 802.3ad ael
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Displaying LACP LAG Information

```
admin@XorPlus# run show interface aggregate-ethernet ael
Physical interface: ael, Enabled, Physical link is Down
Interface index: 53
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 60:eb:69:d2:9c:d7, Hardware address: 60:eb:69:d2:9c:d7
Traffic statistics:
5 sec input rate 0 bits/sec, 0 packets/sec
5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....0
Output Packets.....0
Input Octets.....0
Output Octets.....0
Aggregated link protocol: LACP
Minimum number of selected ports: 4
Members Status Port Speed
-----
ge-1/1/1 up(active) Auto
ge-1/1/2 up(active) Auto
ge-1/1/3 up(active) Auto
ge-1/1/4 up(active) Auto
```

Storm Control in Ethernet Port Configuration

You can configure unicast, multicast, and broadcast storm control in packets per second, kilobits per second or percentage of physical link speed. The storm control function can permit the max rate of unicast, multicast, and broadcast traffics on ingress port.



- Storm-control function is invalid for the known unicast traffic. It works on the unknown unicast, all multicast and broadcast traffic.
- pps, kbps and ratio are mutually exclusive and can not be configured at the same time.

Configuring Storm Control in packet per second on ingress port

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control broadcast pps 10000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control multicast pps 10000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control unicast pps 10000
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Storm Control in ratio on ingress port

The ratio means the percentage of the physical link speed.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 storm-control broadcast ratio 10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 storm-control multicast ratio 20
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 storm-control unicast ratio 30
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Storm Control in kilobits per second on ingress port

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control broadcast kbps 1000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control multicast kbps 1000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 storm-control unicast kbps 1000
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

NOTE:

For storm control in kbps, the actual value of storm control and the configured value are not exactly the same, the correspondence is:

When the configured value is between $64*n$ and $64*n+63$, where n is an integer, the actual value of the storm control is $64*n$. For example, when the configured value is 200 kbps (between $64*3$ and 255), the actual value of the storm control is $64*3=192$ kbps. There is a special case, when the configured value is between 0 and 63, the actual value of the flow control is 64 kbps.

UDLD Configuration

UDLD (Unidirectional Link Detection) using for detecting optical fiber unidirectional link. It supports two modes of operation: normal mode(the default) and aggressive mode. In normal mode, UDLD can detect unidirectional links due to mis-connected interfaces. In aggressive mode, UDLD can also detect unidirectional links due to one-way traffic, twisted-pair links, and mis-connected interfaces. You can enable UDLD globally or on specific ports. When UDLD detects uni-directional fault, the port status will be setted Disabled(UDLD) and down, and there is no udld neighbor. In addition, UDLD can detect the self loop port and disable this port. Self loop means TX links to RX on the same port.

Configure UDLD normal mode on global

```
admin@XorPlus# set protocols udld disable false
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure UDLD aggressive mode on global

```
admin@XorPlus# set protocols udld disable false
admin@XorPlus# set protocols udld aggressive true
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure UDLD normal mode on Specific Port

```
admin@XorPlus# set protocols udld interface te-1/1/27 disable false
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure UDLD aggressive mode on Specific Port

```
admin@XorPlus# set protocols udld interface te-1/1/27 disable false
admin@XorPlus# set protocols udld interface te-1/1/27 aggressive true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure UDLD Message-interval

```
admin@XorPlus# set protocols udld message-interval 20
admin@XorPlus# commit
Commit OK.
Save done.
```

UDLD status as follows when the link is good

```
admin@Xorplus# run show udld neighbors
  Port          Device Name      Device ID          Port ID          state
  -----
te-1/1/27      3295             QTFQXI10700003   te-1/1/49       bi-directional

admin@Xorplus# run show udld interface te-1/1/27
Interface te-1/1/27
```

```

-----
Udld enabled, aggressive mode
Current bidirectional state: bi-directional
Current phase: advertisement
Message interval: 15s
Timeout interval: 5s
  neighbor 1
  -----
  Expiration time: 43.58s
  Device ID: QTFQXI10700003
  Port ID: te-1/1/49
  Message interval: 20s
  Timeout interval: 5s
  Device name: 3295

```

UDLD status as follows when UDLD detect uni-directional fault

```

admin@Xorplus# run show udld neighbors
  Port          Device Name      Device ID        Port ID         state
  -----
admin@Xorplus# run show udld interface te-1/1/27
Interface te-1/1/27
-----
Udld enabled
Current bidirectional state: uni-directional
Current phase: linkdown
Message interval: 7s
Timeout interval: 5s

admin@XorPlus# run show interface gigabit-ethernet te-1/1/27
Physical interface: te-1/1/27, Disabled(UDLD), error-discard False, Physical link is Down
Interface index: 27, SFP type: SR/850
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit  ingress:0, egress:0
Current address: 00:90:4c:06:a5:73, Hardware address: 00:90:4c:06:a5:73
Traffic statistics:
  5 sec input rate 0 bits/sec, 0 packets/sec
  5 sec output rate 0 bits/sec, 0 packets/sec
  Input Packets.....16
  Output Packets.....1360
  Input Octets.....1797
  Output Octets.....157178

```

UDLD status when self loop link

The example is that tx connect with rx on the same port te-1/1/27. The udld status will be transmit-to-receive loop, and the port status will be setted Disabled(UDLD) and down, and there is no udld neighbor.

```

admin@Xorplus# run show udld neighbors
  Port          Device Name      Device ID        Port ID         state
  -----
admin@Xorplus# run show udld interface te-1/1/27
Interface te-1/1/27
-----

```

```

Ulld enabled, aggressive mode
Current bidirectional state: transmit-to-receive loop
Current phase: linkdown
Message interval: 7s
Timeout interval: 5s

admin@XorPlus# run show interface gigabit-ethernet te-1/1/27
Physical interface: te-1/1/27, Disabled(UDLD), error-discard False, Physical link is Down
Interface index: 27, SFP type: SR/850
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit  ingress:0, egress:0
Current address: 00:90:4c:06:a5:73, Hardware address: 00:90:4c:06:a5:73
Traffic statistics:
 5 sec input rate 0 bits/sec, 0 packets/sec
 5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....16
Output Packets.....1360
Input Octets.....1797
Output Octets.....157178

```

Recover port status Disabled(UDLD)

You can through down, up the port to recover the status Disabled(UDLD)

```

admin@XorPlus# set interface gigabit-ethernet te-1/1/27 disable true
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet te-1/1/27 disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet te-1/1/27
Physical interface: te-1/1/27, Enabled, error-discard False, Physical link is Down
Interface index: 27, SFP type: SR/850
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit  ingress:0, egress:0
Current address: 00:90:4c:06:a5:73, Hardware address: 00:90:4c:06:a5:73
Traffic statistics:
 5 sec input rate 0 bits/sec, 0 packets/sec
 5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....16
Output Packets.....1360
Input Octets.....1797
Output Octets.....157178

```

Or you can using follow command to recover the status Disabled(UDLD)

```

admin@XorPlus# run clear ulld
admin@XorPlus# run show interface gigabit-ethernet te-1/1/27
Physical interface: te-1/1/27, Enabled, error-discard False, Physical link is Down
Interface index: 27, SFP type: SR/850
Description:
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Interface rate limit  ingress:0, egress:0

```

```

Current address: 00:90:4c:06:a5:73, Hardware address: 00:90:4c:06:a5:73
Traffic statistics:
 5 sec input rate 0 bits/sec, 0 packets/sec
 5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....16
Output Packets.....1360
Input Octets.....1797
Output Octets.....157178

```

Port Security Configuration

Port security is a layer two traffic control feature on Pica8 switches. It enables an administrator to configure individual switch ports to allow only a specified number of source MAC addresses to ingress through the port. Port security enables the switch administrator to prevent unauthorized devices from gaining access to the network. Port security is normally enabled on access layer switches for this purpose.

Enabling Port Security

Port security is not enabled by default. It can be enabled with default parameters by issuing a single command on an interface:

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-limit 10
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring the Maximum Number of Secure Dynamically Learned MAC Addresses

User can use port security with dynamically learned MAC addresses to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port.

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-limit 5
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show port-security address
Secure Mac Address Table
-----
Vlan MAC Address Type Interfaces
-----
1 00:00:11:11:11:11 dynamic ge-1/1/1
1 00:00:11:11:11:12 dynamic ge-1/1/1
1 00:00:11:11:11:13 dynamic ge-1/1/1
1 00:00:11:11:11:14 dynamic ge-1/1/1
1 00:00:11:11:11:15 dynamic ge-1/1/1

```

```

-----
MAC age time :300s
admin@XorPlus#

```

Configuring Static Secure MAC Addresses on a Port

User can use port security with static MAC addresses to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port.

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-address
00:00:23:23:23:23 vlan 1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-address
00:00:23:23:23:24 vlan 1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-address
00:00:23:23:23:25 vlan 1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-address
00:00:23:23:23:26 vlan 1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security mac-address
00:00:23:23:23:27 vlan 1
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show port-security address
Secure Mac Address Table
-----
Vlan MAC Address Type Interfaces
-----
1 00:00:23:23:23:23 static ge-1/1/1
1 00:00:23:23:23:24 static ge-1/1/1
1 00:00:23:23:23:25 static ge-1/1/1
1 00:00:23:23:23:26 static ge-1/1/1
1 00:00:23:23:23:27 static ge-1/1/1
-----
MAC age time :300s
admin@XorPlus#

```

Configuring Port Security with Sticky MAC Addresses on a Port

Port security with sticky MAC addresses retains dynamically learned MAC addresses when the link is down and restores the MAC addresses when the link is up.

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security sticky true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Secure MAC Address Aging Time

The aging time is global whether port security is configured or not.

```

admin@XorPlus# set interface ethernet-switching-options mac-table-aging-time 100
admin@XorPlus# commit

```

```
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Port Security Violation Mode on a Port

Port security can be configured to take one of four actions upon detecting a violation:

- **protect** (default) - Frames from MAC. Addresses other than the allowed addresses are dropped. Traffic from allowed addresses is permitted to pass normally.
- **restrict** - Like protect mode but generates a syslog message and increases the violation counter.
- **shutdown** - The interface is placed into the error-discard state, blocking all traffic.
- **shutdown-temp** - The interface is placed into the error-discard state and blocking all traffic temporarily. After 20 seconds (default), the interface comes up.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security violation ?
Possible completions:
protect Drop packets with unknown source addresses
restrict Drop packets with unknown source addresses and log violation
shutdown Disable interface
shutdown-temp Disable interface temporarily(20 seconds for the default)
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security violation restrict
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Port Security Auto-recovery Time

When the port security violation mode is configured to shutdown-temp, user can configure the recovery interval with the command below.

```
admin@XorPlus# set interface ethernet-switching-options port-error-discard timeout 30
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Recovering the Port in Error-discard

When the port security violation mode is configured to shutdown, the port will be placed into the error-discard state after detecting a violation. User can recover the port with the following command.

```
admin@XorPlus# run clear port-security port-error
Clear done.
admin@XorPlus#
```

Configuring Port Security Block Mode on a Port

Port security can be configured to take one of five block actions:

- **all** – All traffic are not permitted to forward normally on egress.
- **broadcast**– Broadcast packets will be blocked on egress, but unknown uni/multi cast addresses can forwards normally.
- **multicast** – Only the multicast packets will be dropped.
- **uni-multi-cast**- The unknown uni/multi cast packets will be blocked on egress.
- **unicast** - Only the unknown unicast packets will be dropped.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security block ?
Possible completions:
all Block broadcast and unknow addresses
broadcast Block broadcast address
multicast Block unknow multicast addresses
uni-multi-cast Block unknow uni/multi cast addresses
unicast Block unknow unicast addresses
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 port-security block broadcast
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Displaying Port Security Settings

To display port security settings, enter this command:

```
admin@XorPlus# run show port-security address
Secure Mac Address Table
-----
Vlan MAC Address Type Interfaces
-----
1 00:00:11:11:11:11 dynamic ge-1/1/1
1 00:00:11:11:11:12 dynamic ge-1/1/1
1 00:00:11:11:11:13 dynamic ge-1/1/1
1 00:00:11:11:11:14 dynamic ge-1/1/1
1 00:00:11:11:11:15 dynamic ge-1/1/1
1 00:00:23:23:23:23 static ge-1/1/1
1 00:00:23:23:23:24 static ge-1/1/1
1 00:00:23:23:23:25 static ge-1/1/1
1 00:00:23:23:23:26 static ge-1/1/1
1 00:00:23:23:23:27 static ge-1/1/1
-----
MAC age time :100s
admin@XorPlus# run show port-security brief
System MAC limit : 32767
Secure port DynamicMacLim CurrentAddr ViolationCount Action
-----
ge-1/1/1 5 10 213940 restrict
-----
admin@XorPlus#
admin@XorPlus# run show port-security interface gigabit-ethernet ge-1/1/1
Interface ge-1/1/1
-----
Port security : enabled
Violation action : restrict
```

```

Block type : broadcast
Sticky : true
Dynamic MAC limit : 5
Total MAC addresses : 10
Configured MAC addresses : 5
Sticky MAC addresses : 5
Security violation count : 286062
admin@XorPlus#

```

Disabling Port Security

To disable port security, enter this command:

```

admin@XorPlus# delete interface gigabit-ethernet ge-1/1/1 port-security
Deleting:
port-security {
mac-limit: 5
violation: "restrict"
mac-address 00:00:23:23:23:23 {
vlan 1 {
}
}
mac-address 00:00:23:23:23:24 {
vlan 1 {
}
}
mac-address 00:00:23:23:23:25 {
vlan 1 {
}
}
mac-address 00:00:23:23:23:26 {
vlan 1 {
}
}
mac-address 00:00:23:23:23:27 {
vlan 1 {
}
}
sticky: true
block: "broadcast"
}
OK
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

802.1x Configuration

IEEE 802.1x is a Port-based Network Access Control Protocol, which provides authentication mechanism for both LAN and WLAN devices. It allows user to use LAN/WAN ports only after successful authentication. However, if authentication fails, user is not able to use LAN/WLAN ports, even if there are physical or wireless connections. 802.1x framework consists of three elements: a supplicant, an authenticator, and an authentication server.



In the 2.6.4 previous version, Dot1x only support EAP-MD5 authentication method
In 2.6.4 version or later, Dot1x support EAP-MD5,EAP-PEAP,EAP-TLS,EAP-TTLS authentication methods.

Enabling/ disabling 802.1x on the Interface

Enable 802.1x on an access port:

802.1x is not enabled by default. It can be enabled by issuing a single command on an interface:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 disable false
admin@XorPlus#commit
```

Disable 802.1x on an access port:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 disable true
admin@XorPlus#commit
```

Configuring 802.1x Port Control Mode of Interface

Configuring auto mode:

By default, the port-control mode is “auto”.

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 port-control auto
ERROR: The same value is set to node"protocols dot1x interface ge-1/1/8 port-control".
admin@XorPlus#commit
```

Configuring force-authorized mode:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 port-control force-authorized
admin@XorPlus#commit
```

Configuring force-unauthorized mode:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 port-control force-unauthorized
admin@XorPlus#commit
```

Configuring 802.1x Radius Server

Configuring radius server and the shared-key.

```
admin@XorPlus# set protocols dot1x aaa radius authentication server-ip 1.1.5.1 shared-key
123
admin@XorPlus#commit
```

Configuring the Value of Reauth-period

By default, the value of reauth-period is 3600s.

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 reauth-period 600
admin@XorPlus# commit
```

Enable/disable Re-Authentication

By default, the re-authentication is enabled.

Enable re-authentication:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 reauthentication disable false
admin@XorPlus# commit
```

Disable re-authentication:

```
admin@XorPlus# set protocols dot1x interface ge-1/1/8 reauthentication disable true
admin@XorPlus# commit
```

Displaying 802.1x Settings on Auto Mode

```
admin@XorPlus# run show dot1x interface
Interface Client Status
-----
ge-1/1/8 08:9e:01:39:1a:fe AUTHORIZED

admin@XorPlus# run show dot1x interface gigabit-ethernet ge-1/1/8
Dot1x Info for ge-1/1/8
-----
PortEnabled = true
PortControl = AUTO
QuietPeriod = 60
ServerTimeout = 30
ReAuthentication = true
ReAuthPeriod = 3600
Dot1x Authenticator Client
-----
Supplicant = 08:9e:01:39:1a:fe
Port Status = AUTHORIZED
Auth SM State = AUTHENTICATED
Auth BEND SM State = IDLE
```

Deleting Protocols of dot1x

```
admin@XorPlus# delete protocols dot1x
Deleting:
  dot1x {
    interface "ge-1/1/8" {
      disable: false
      reauthentication {
        disable: false
      }
      reauth-period: 600
    }
  }
  aaa {
    radius {
      authentication {
        server-ip 1.1.5.1 {
```

```

        shared-key: "123"
    }
}
}
admin@XorPlus#commit

```

Buffer management

 Buffer management is not supported on Kontron MSH8920.

Introduction

When the network congestion, the port don't send data immediately, to prevent data loss, device can transmit the temporay storage of data to the data buffer.

1. Accounting Resources

The chip has the fixed-sized cells. Each cell is 208 bytes long. The first cell used by a packet contains two payloads. The first payload is a 64-byte packet descriptor, and the second payload contains 144 bytes of packet data. Every cell, after the first one used by a packet, uses the complete 208 bytes for packet data . A packet use at least one cell. The length of packet decides how many cell the packet uses. For example untagged packeets 64 bytes to 144 bytes long use one cell. A 352-byte packet would require two cells. And packets from 1393-bytes to 1600-bytes long use eight cells.

2. Memory Spaces

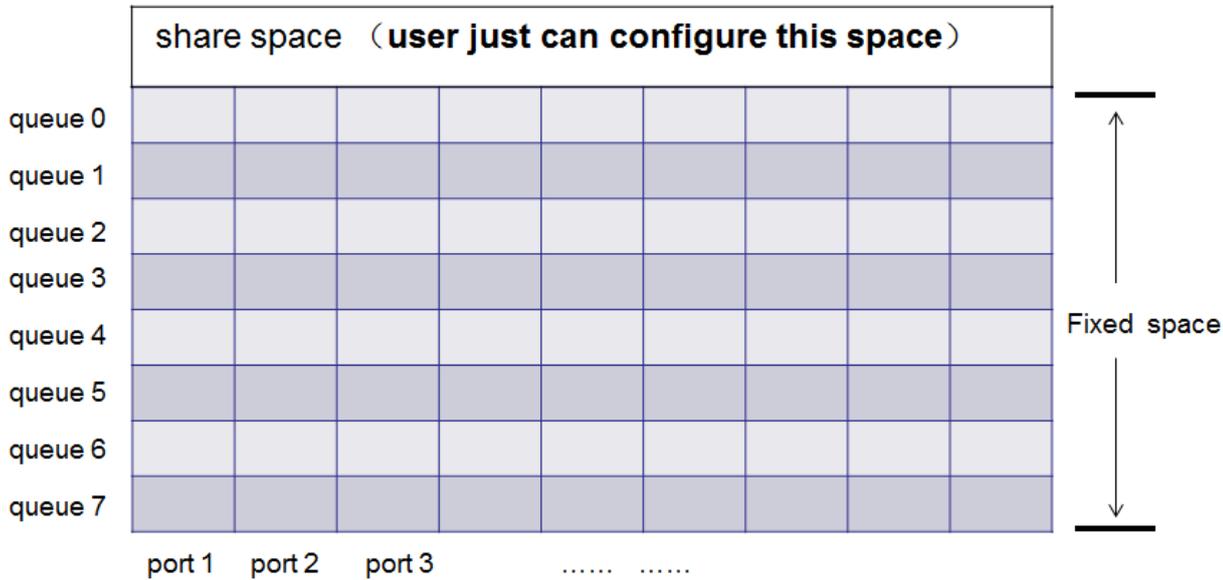
The cell data buffer is divided into fixed and shared spaces.

- Fixed space

Fixed space provide minimum guaranteed space on a per priority group and per port basis. This space cannot be shared with other ports. When the port congestion occurs, it will occupy fixed space first, shared space will be used when the fixed space is full. The Fixed space should not be too big, otherwise it will waste space resources.

- Shared space

The shared space will be used by all ports and priority groups when the fixed space run out. When the shared space is used up, the packet is discarded. If a queue don't store packet, the others queue will grab this queue space. For a queue,all of ports transmit packets will first come first served, if the used resource reach the threshold, packets will be dropped after arrived.The shared space can be configured dynamic or static threshold at the logical level of the queue.



Command configuration

The device provide separate UC (unicast) and MC (multicast) queue. For the unicast queue and multicast queue. The UC queues are used to store known unicast packets and can be flexibly assigned while the MC queues are used to store DLF, multicast, broadcast and mirrored packets. While the CPU port does not support separate UC and MC queues, the CPU port does have a total of 48 queues available, and these queues are eligible for both UC and MC traffic based on CPU Cos mapping. We have different configuration to deal with share space

1.multicast queue

static mode: This is the default configuration

- The value of the static threshold is configured for the multicast queue in the shared space. This avoids the case that the port enabled flow control or PFC affects the forwarding of other ports. But the disadvantage is less efficient use of memory.

dynamic mode: Need to configure

- The memory threshold is adjusted dynamically according to the available shared memory in this mode. The thresholds will reduce when the remaining shared memory reduces. The thresholds will rise up when the remaining shared memory rises up.

For example, enable queue 0 dynamic-shared

```
root@XorPlus# set interface ethernet-switching-options buffer egress-queue 0
mc-queue-dynamic-shared true
root@XorPlus# commit
```

2.unicast queue

dynamic mode

- dynamic threshold is the default mode for the unicast queue and it can't configure static threshold.

3.dynamic threshold ratio configuration

The dynamic mode both of multicast (must enable dynamic threshold) and unicast queue can configure the max dynamic threshold ratio. The threshold is still dynamically adjusted until reaching the max ratio configured.

For example, configure dynamic threshold ratio 50% for queue 0

```
root@XorPlus# set interface ethernet-switching-options buffer egress-queue 0 shared-ratio
50
root@XorPlus# commit
```

Check the configuration result

```
root@XorPlus# run show interface egress-buffer
Cosq Buffer-Shared-Ratio Multicast-Queue-Dnamic-Shared-Buffer
0      50                  Enable
1      33                  Disable
2      33                  Disable
3      33                  Disable
4      33                  Disable
5      33                  Disable
6      33                  Disable
7      33                  Disable
```

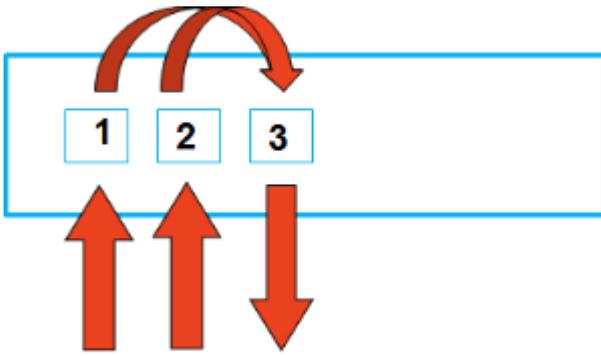
For the maximum shared-area ratio for a queue, the percentage values 0 to 100 are divided into 10 rages.The table below shows the effective values that correspond to the configured values of ratio-value.

Value of <i>ratio-value</i>	Effective value
0 to 1	1
2 to 3	3
4 to 7	6
8 to 16	11
17 to 29	20
30 to 42	33
43 to 60	50
61 to 76	67
77 to 86	80
87 to 100	89

notice

Now,multicast queue just support eight 1~67effective ratio value on all platforms

Application notice



As show above: send packets from port1 and port2 to port3 in 10G speed in queue 0, and port3 will be in congestion.

For the multicast flow, if flow control is enabled on port1 and port2 and configure static threshold, the packets also perhaps will be dropped in port 3 egress queue. Ingress don't send pause frame because egress's shared-ratio is smaller than ingress's and egress reach it's threshold before igress reach.

So if users meet this issue, in this case, we suggest users configure dynamic mode. port1 and port2 will send out pause frame before port3 queue limit is reached, and accept burst packets to avoid packet drop.

Layer 2 Switching Configuration

This chapter describes the configuration steps of Layer 2 switching, including MAC address learning, LLDP, LACP, 802.1Q VLAN, flow control, mirroring, storm control, and the Spanning Tree Protocol (STP/RSTP/MSTP).

- Static MAC entries and Dynamic MAC Address Learning
- Cut-Through Switching Method
- MLAG Configuration Guide
 - Configuring a Basic MLAG Step-by-Step Procedure
 - Configuring a Basic MLAG Example
 - Configuring Switch A with Static and LACP LAG
 - Configuring Switch B with Static and LACP LAG
 - Configuring Switch C with LACP and LAG
 - View the MLAG Internal and Neighbor Status
 - Configuring a MLAG domain with MSTP example
 - Configuring Switch A with LACP LAG
 - Configuring Switch B with LACP LAG
 - Configuring an Aggregation Interface to VLAN Members
 - Configuring Switch C and Switch D with LACP LAG
 - Configure Mlag and Active-Active-Vrrp example
 - Configure on Switch A
 - Configure on Switch B
 - Configure on Switch C
 - Configure on Switch D
 - Configuring a MLAG with vxlan Example
 - 01 Configuring mlag in Switch A and Switch B
 - 02 Configuring Lag in Switch C and Switch D
 - 03 Configuring vxlan in Switch A and Switch B
- Q-in-Q Basic Port Configuration
- MSTP Configuration
- MSTP Configuration Example
- PVST Configuration
- PVST Configuration Example
- Configuring Mirroring

- Configuring Mirroring Guide
- BPDU Tunneling Configuration
- Configuring IPv6 RA Guard
- LFS Configuration
- Symmetric Hash for ECMP and LAG Configuration Example

Static MAC entries and Dynamic MAC Address Learning

You can configure a static MAC entry in the FDB and manage dynamic MAC address learning (for example, configuring aging time or deleting the dynamic MAC address entry).

Configuring a Static MAC Entry and Managing the FDB

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 static-ethernet-switching
mac-address 22:22:22:22:22:22 vlan 1
admin@XorPlus# set interface ethernet-switching-options mac-table-aging-time 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show ethernet-switching table
Total entries in switching table: 2
Static entries in switching table: 0
Dynamic entries in switching table: 2
VLAN MAC address Type Age Interfaces
-----
1 00:22:be:96:f2:83 Dynamic 60 ge-1/1/1
1 00:22:be:96:f2:84 Dynamic 60 ge-1/1/2
admin@XorPlus# run clear ethernet-switching table all
admin@XorPlus# run show ethernet-switching table
Total entries in switching table: 0
Static entries in switching table: 0
Dynamic entries in switching table: 0
VLAN MAC address Type Age Interfaces
-----
admin@XorPlus#
```

Cut-Through Switching Method

By default, the switch forwards the packets in a cut-through switching method. That is, the switch begins forwarding a packet before the entire frame is received, normally as soon as the destination address is processed. This process reduces latency. Error handling is performed by the destination devices. User can configure the switch to the "store-and-forward" method with the commands below.

 4610 series switches don't support cut-through mode.

Configuring your Switch to the Store-and-Forward Method

```
admin@XorPlus# set interface cut_through_mode false
admin@XorPlus# commit
```

```
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

MLAG Configuration Guide

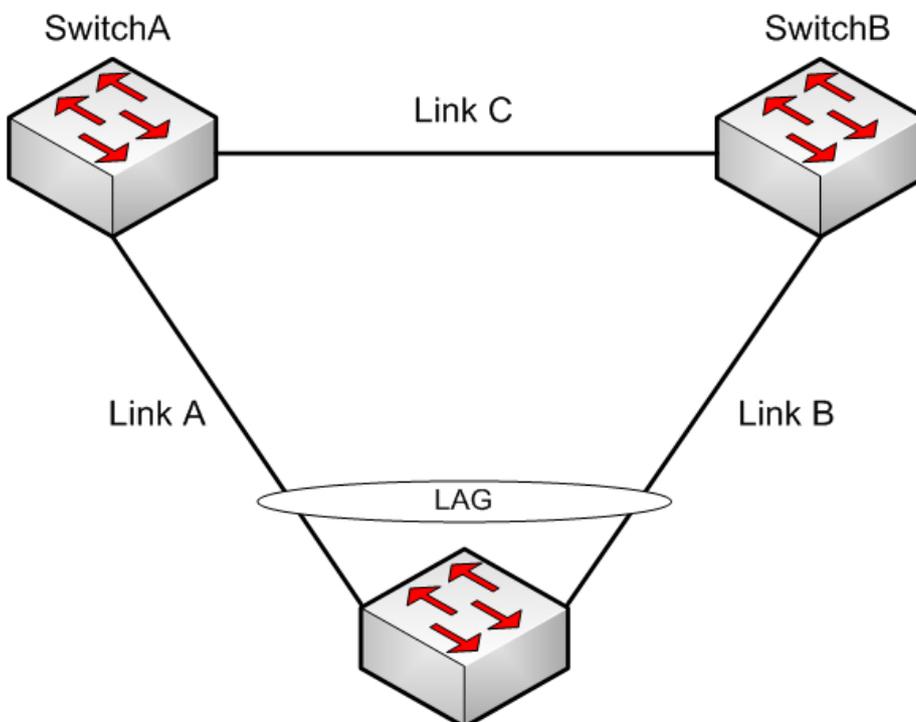
Traditionally, an aggregation interface is a logical interface that is used to increase bandwidth or availability by using more than one physical interface in a switch. Multi-chassis LAG (MLAG) can form a logical aggregation interface to multiple switches.

In Figure 1-1, switch A and C are connected by link A; switch B and C are connected by link B. In switch C, link A and B have formed an aggregation interface to balance traffic. Meanwhile, switches A and B have formed an MLAG using link A and B. For communication such as MAC entries between the members of the MLAG are learned by the MLAG and must be synchronized. In Figure 1, synchronization between switch A and B is achieved via link C which is used to connect switch A and B as the channel interface. The number of links which connect switch A and C, or B and C, cannot be more than 1.

Important Things To Know About MLAG

There are two issues in the MLAG: MAC entry synchronization and broadcast traffic control. MAC entry synchronization means that the MAC entry learned by the interface must be synchronized by the peer switch. In the current version, only 2 nodes in a MLAG and using L2 traffic to communicate between notes is supported.

Figure 1-1



Configuring MLAG domain-id

The *domain-id* command assigns an MLAG ID to an aggregation interface. MLAG neighbor switches from an MLAG when each switch configures the same MLAG-ID to an aggregation interface. Only one MLAG domain ID can be assigned to an aggregation interface. MLAG domain ID's are unique. That is to say, an MLAG domain-ID cannot be assigned to more than one aggregation interface.

```
admin@XorPlus# set interface aggregate-ethernet ae22 aggregated-ether-options mlag
domain-id 1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring MLAG Peer

The *peer* command specifies the neighbor's IP address for a MLAG domain. The synchronization messages between MLAG peers are sent to the neighbor's IP address.

```
admin@XorPlus# set interface aggregate-ethernet ae22 aggregated-ether-options mlag peer
10.0.0.1 peer-link ae24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring MLAG Priority

The priority uses master/slave negotiation between the two neighbor switches. The *priority* command assigns a MLAG priority.

```
admin@XorPlus# set interface aggregate-ethernet ae22 aggregated-ether-options mlag priority
4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring MLAG hello-interval

The *hello-interval* command configures the hello message interval in both directions between MLAG neighbors. If the neighbor switch is pinged four times and a reply is not received during the interval, the MLAG neighbor switches revert to their independent state.

```

admin@XorPlus# set interface aggregate-ethernet ae22 aggregated-ether-options mlag
hello-interval 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

configuration Consistency Check

From version 2.9.1, add the function of configuration consistency check between two peer nodes. Mac address sync can work well based on the same MLAG interface configuration between two nodes.

Configuration consistency fields on MLAG interface as below:

1. Vlan id (native vlan id and vlan member)
2. Tagged / untagged (vlan member is tagged or untagged)
3. Mac-learning (true / false)

Match or mismatch result is shown in the command "show mlag internal" .

```

admin@XorPlus# run show mlag internal
Domain-id  Local-LAG  Flood  MAC-sync  State      Conf-Consistency  Role
-----
2           ae1         false true      FULL      CONFIG_MATCH      MASTER
1           ae2         false true      FULL      CONFIG_MATCH      MASTER

```

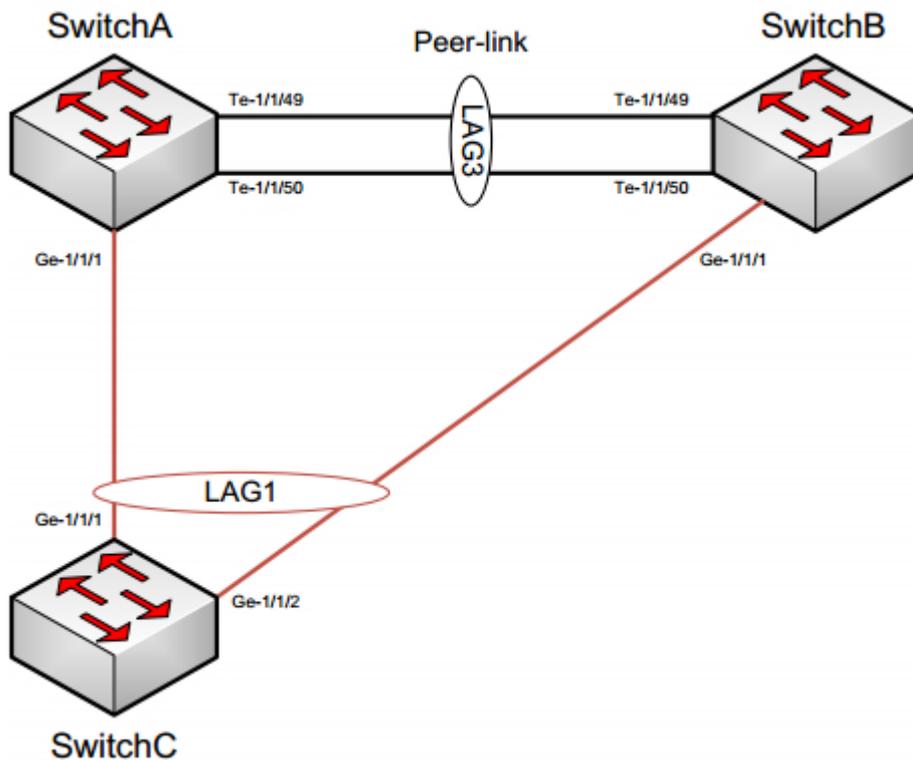


We don't recommend to configure system-id specially for MLAG. MLAG can automatically use mac-address of the in-band interface as system-id.

- [Configuring a Basic MLAG Step-by-Step Procedure](#)
- [Configuring a Basic MLAG Example](#)
 - [Configuring Switch A with Static and LACP LAG](#)
 - [Configuring Switch B with Static and LACP LAG](#)
 - [Configuring Switch C with LACP and LAG](#)
 - [View the MLAG Internal and Neighbor Status](#)
- [Configuring a MLAG domain with MSTP example](#)
 - [Configuring Switch A with LACP LAG](#)
 - [Configuring Switch B with LACP LAG](#)
 - [Configuring an Aggregation Interface to VLAN Members](#)
 - [Configuring Switch C and Switch D with LACP LAG](#)
- [Configure Mlag and Active-Active-Vrrp example](#)
 - [Configure on Switch A](#)
 - [Configure on Switch B](#)
 - [Configure on Switch C](#)
 - [Configure on Switch D](#)

- Configuring a MLAG with vxlan Example
 - 01 Configuring mlag in Switch A and Switch B
 - Configuring Lag in Switch C and Switch D
 - 02 Configuring Lag in Switch C and Switch D
 - 03 Configuring vxlan in Switch A and Switch B

Configuring a Basic MLAG Step-by-Step Procedure



1. Configure the number of LAGs on Switch C. Add member interfaces to the aggregated ethernet interface on Switch C.
2. Configure the number of MLAG member LAGs on both Switch A and Switch B. Add member interfaces to the aggregated ethernet interface on both Switch A and Switch B.
3. Configure the number of MLAG peer-link LAGs on both Switch A and Switch B. Add member interfaces to the aggregated ethernet interface on both Switch A and Switch B.
4. Configure the L3 interface IP address on both Switch A and Switch B for peer-to-peer communication.
5. Configure the same domain-id number on both MLAG peers on Switch A and Switch B.
6. Configure different system-id's on both MLAG peers on Switch A and Switch B.
7. Configure the peer IP address for MLAG peer connect on both Switch A and Switch B.
8. Configure the LAGs for MLAG peer-link connects on both Switch A and Switch B.

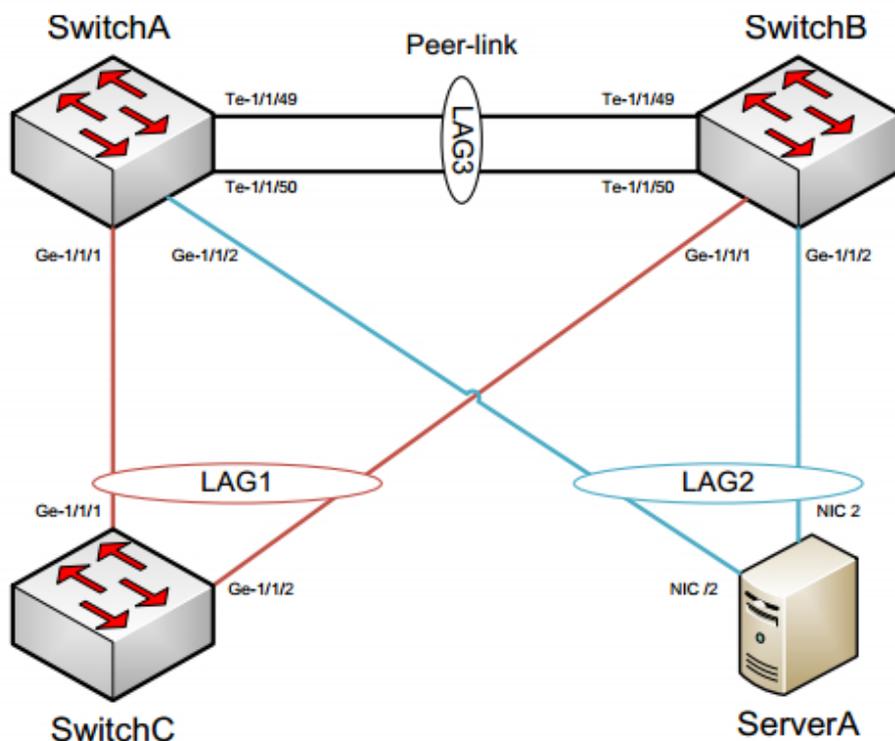
Configuring a Basic MLAG Example

Figure 2 illustrates MLAG configured between Switch A and Switch B, the MLAG connections between the neighboring switches, and two Network Devices.

The MLAG switches connect through a LACP LAG to Switch C.

The MLAG switches connect through a static LAG to Server A.

Figure 2



- Configuring Switch A with Static and LACP LAG
- Configuring Switch B with Static and LACP LAG
- Configuring Switch C with LACP and LAG
- View the MLAG Internal and Neighbor Status

Configuring Switch A with Static and LACP LAG

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lACP enable
true
admin@XorPlus# set interface aggregate-ethernet ae2
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.
Save done.
admin@XorPlus#
```

Configuring an Aggregation Interface to VLAN Members

```
admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 l3-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching
native-vlan-id 4094
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 Interface IP Address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.1 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the domain-id for the MLAG

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id
1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

```

admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring the Peer IP Address and the Peer-link for the MLAG Domain Peer

```

admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer
10.10.0.2 peer-link ae3
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer
10.10.0.2 peer-link ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B with Static and LACP LAG

```

admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae2
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring an Aggregation Interface to VLAN Members

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 l3-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode

```

```
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching
native-vlan-id 4094
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 Interface IP Address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.2 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the domain-id for the MLAG Domain

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id
1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

Configuring the Peer IP Address and the peer-link for the MLAG Domain Peer

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer
10.10.0.1 peer-link ae3
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer
10.10.0.1 peer-link ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch C with LACP and LAG

```

admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configuring an Aggregation Interface to VLAN Members

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

View the MLAG Internal and Neighbor Status

View the MLAG Internal and Neighbor Status of Switch A

```

admin@XorPlus# run show mlag internal
Domain-id  Local-LAG  Flood  MAC-sync  State        Conf-Consistency  Role
-----
2          ae1        false true      FULL         CONFIG_MATCH      MASTER
1          ae2        false true      FULL         CONFIG_MATCH      MASTER

admin@XorPlus# run show mlag peer 1
Peer          System-id          State  Link-status
-----
10.10.0.2     c8:0a:a9:9e:14:a4  FULL  UP

admin@XorPlus# run show mlag peer 2
Peer          System-id          State  Link-status
-----
10.10.0.2     c8:0a:a9:9e:14:a4  FULL  UP

```

View the MLAG Internal and Neighbor Status of Switch B

```
admin@XorPlus# run show mlag internal
Domain-id  Local-LAG  Flood  MAC-sync  State        Conf-Consistency  Role
-----
2          ae1        false true      FULL         CONFIG_MATCH      SLAVE
1          ae2        false true      FULL         CONFIG_MATCH      SLAVE

admin@XorPlus# run show mlag peer 1
Peer          System-id          State  Link-status
-----
10.10.0.1     e8:9a:8f:50:3d:30  FULL  UP

admin@XorPlus# run show mlag peer 2
Peer          System-id          State  Link-status
-----
10.10.0.1     e8:9a:8f:50:3d:30  FULL  UP
```

If the configuration is inconsistent, the mlag internal as follows:

```
admin@XorPlus# run show mlag internal
Domain-id  Local-LAG  Flood  MAC-sync  State        Conf-Consistency  Role
-----
1          ae2        false true      FULL         CONFIG_MISMATCH   MASTER

admin@XorPlus# run show mlag internal 1 detail
Local-LAG: ae2
Native VlanID: 1941
Mac Learning: true
Untagged VlanID:
Tagged VlanID:
157, 203
```

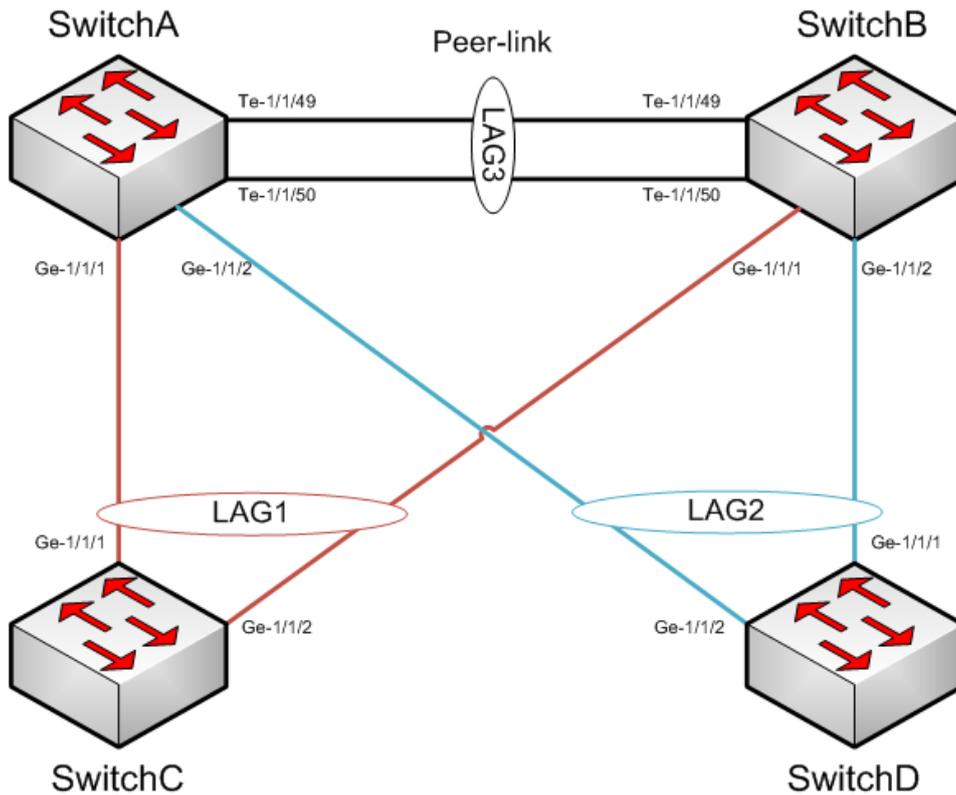
Configuring a MLAG domain with MSTP example

Figure 3 illustrates MLAG configured between Switch A and Switch B; the MLAG connections between the neighboring switches as well as two Network Devices.

The MLAG switches connect through a LACP LAG to Switch C.

The MLAG switches connect through a LACP LAG to Switch D.

Figure 3



- Configuring Switch A with LACP LAG
- Configuring Switch B with LACP LAG
- Configuring an Aggregation Interface to VLAN Members
- Configuring Switch C and Switch D with LACP LAG

Configuring Switch A with LACP LAG

```

admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lACP enable
true
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options lACP enable
true
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring an Aggregation Interface to VLAN Members

```
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 l3-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 interface IP address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.1 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the domain-id for the MLAG

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id
1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the peer IP address and the peer-link for the MLAG domain peer

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer
10.10.0.2 peer-link ae3
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer
10.10.0.2 peer-link ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B with LACP LAG

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring an Aggregation Interface to VLAN Members

```
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 l3-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching
```

```
native-vlan-id 4094
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 interface IP address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.2 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the domain-id for the MLAG domain

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id
1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the peer IP address and the peer-link for the MLAG domain peer

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer
10.10.0.1 peer-link ae3
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer
10.10.0.1 peer-link ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch C and Switch D with LACP LAG

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# commit
Waiting for merging configuration.
```

```

Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C and Switch D an aggregation interface add to VLAN Members

```

admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members 15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members 16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

View the MLAG internal and neighbor status of Switch A

```

SwitchA# run show mlag internal
Domain-id Local-LAG Flood MAC-sync State Role
-----
2 ae1 false true FULL MASTER
1 ae2 false true FULL MASTER
SwitchA# run show mlag peer 1
Peer System-id State Link-status
-----
10.10.0.2c8:0a:a9:9e:14:a4 FULL UP
SwitchA# run show mlag peer 2
Peer System-id State Link-status
-----
10.10.0.2c8:0a:a9:9e:14:a4 FULL UP
admin@XorPlus#

```

View the MSTP status of Switch A

```

SwitchA# run show spanning-tree mstp interface
MSTP Spanning Tree Interface Status for instance 0
Interface Port ID Designated Designated Bridge Ext Port Int Port State Role
Port ID ID Cost Cost
-----
-----
ae1 128.53 128.53 0.e8:9a:8f:50:3d:30 20000 20000 FORWARDING DESIGNATED

```

```
ae2 128.54 128.54 0.e8:9a:8f:50:3d:30 20000 20000 FORWARDING DESIGNATED
ae3 128.55 128.55 0.e8:9a:8f:50:3d:30 2000 2000 FORWARDING EDGE
```

View the MLAG internal and neighbor status of Switch B

```
SwitchB# run show mlag internal
Domain-id Local-LAG Flood MAC-sync State Role
```

```
-----
2 ae1 false true FULL MASTER
1 ae2 false true FULL MASTER
```

```
SwitchB# run show mlag peer 1
Peer System-id State Link-status
```

```
-----
10.10.0.1e8:9a:8f:50:3d:30 FULL UP
```

```
SwitchB# run show mlag peer 2
```

```
Peer System-id State Link-status
```

```
-----
10.10.0.1e8:9a:8f:50:3d:30 FULL UP
```

```
admin@XorPlus#
```

View the MSTP status of Switch B

```
SwitchA# run show spanning-tree mstp interface
MSTP Spanning Tree Interface Status for instance 0
Interface Port ID Designated Designated Bridge Ext Port Int Port State Role
Port ID ID Cost Cost
```

```
-----
ae1 128.53 128.53 0.e8:9a:8f:50:3d:30 20000 20000 FORWARDING DESIGNATED
```

```
ae2 128.54 128.54 0.e8:9a:8f:50:3d:30 20000 20000 FORWARDING DESIGNATED
```

```
ae3 128.55 128.55 0.e8:9a:8f:50:3d:30 2000 2000 FORWARDING EDGE
```

Configure Mlag and Active-Active-Vrrp example

Requirements

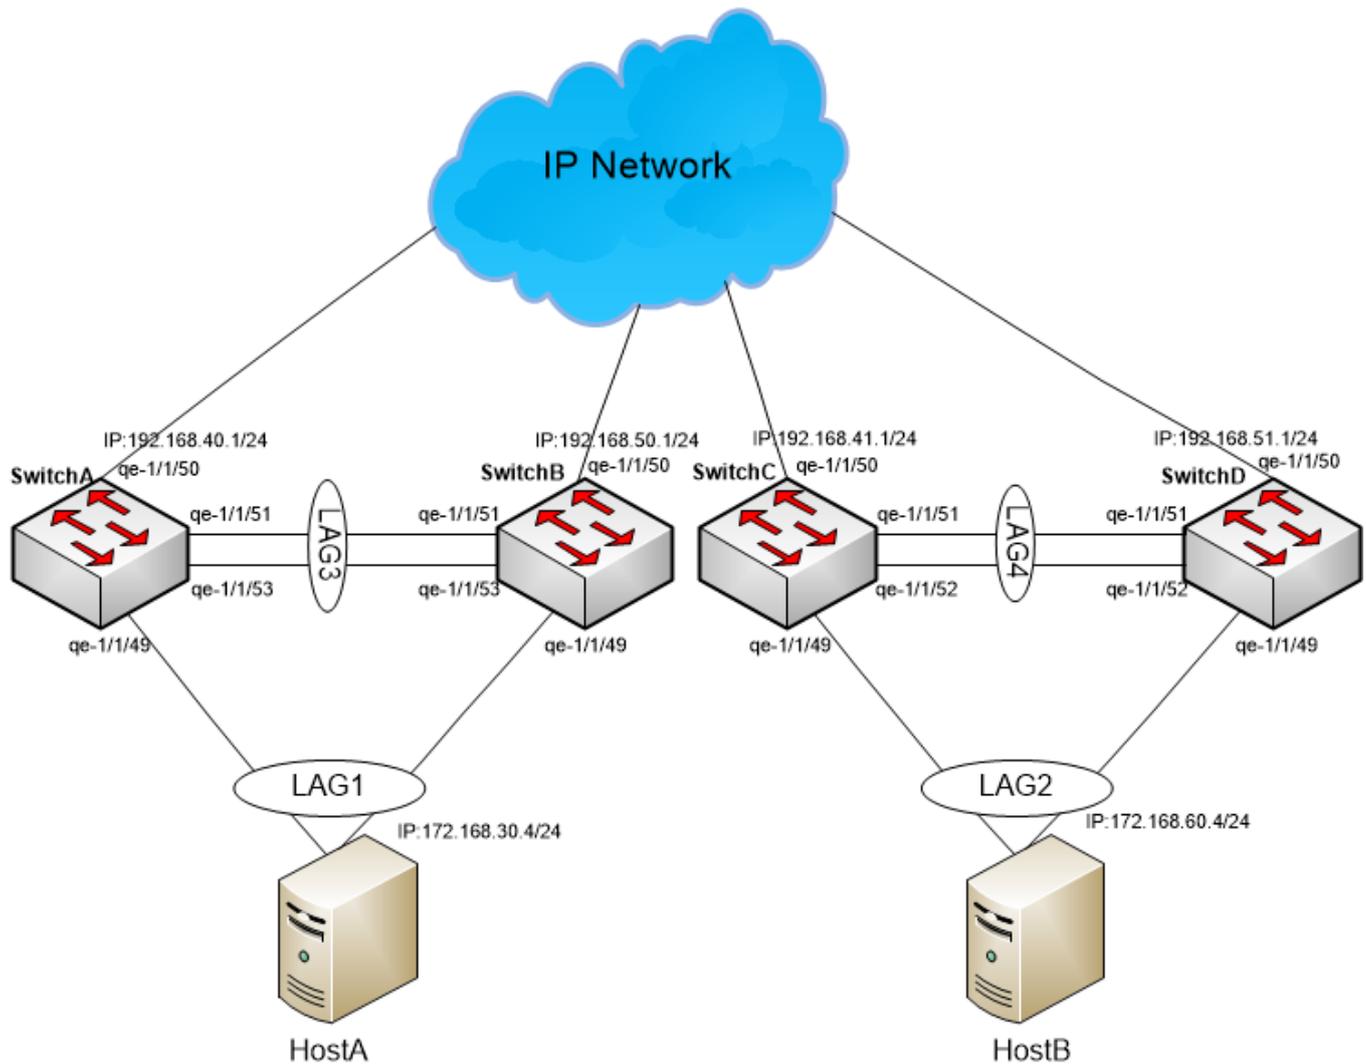
This example uses the following hardware and software components:

- PICOS OS Release, from 2.8.1 version supported Mlag+Vrrp

Overview

In this example, there are two groups mlags and vrrps. It implements the communication between HostA to HostB.

Topology



The topology illustrates two groups MLAGs and Vrrp. One group contains Switch A, Switch B and HostA, another group contains Switch C, Switch D and HostB. This two groups of switches communicate with each other through IP Network.

The two groups of MLAG switches connect through a LACP(Link Aggregation Control Protocol) to HostA and HostB.

- Configure on Switch A
- Configure on Switch B
- Configure on Switch C
- Configure on Switch D

Configure on Switch A

Configure Mlag on Switch A

```

set interface aggregate-ethernet ae1 aggregated-ether-options lACP enable true
set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id 2
set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer 192.168.30.2
peer-link ae3
set interface aggregate-ethernet ae1 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae1 family ethernet-switching vlan members 100
set interface aggregate-ethernet ae3 family ethernet-switching native-vlan-id 300
set interface aggregate-ethernet ae3 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae3 family ethernet-switching vlan members 100
set interface gigabit-ethernet qe-1/1/49 ether-options 802.3ad ae1
set interface gigabit-ethernet qe-1/1/50 family ethernet-switching native-vlan-id 400
set interface gigabit-ethernet qe-1/1/51 ether-options 802.3ad ae3
set interface gigabit-ethernet qe-1/1/53 ether-options 802.3ad ae3
set protocols spanning-tree enable false
set protocols static route 172.168.60.0/24 next-hop 192.168.40.2
set vlan-interface interface vlan-300 vif vlan-300 address 192.168.30.1 prefix-length 24
set vlan-interface interface vlan-100 vif vlan-100 address 172.168.30.2 prefix-length 24
set vlan-interface interface vlan-400 vif vlan-400 address 192.168.40.1 prefix-length 24
set vlans vlan-id 100 l3-interface vlan-100
set vlans vlan-id 300 l3-interface vlan-300
set vlans vlan-id 400 l3-interface vlan-400

```

Configure Vrrp on Switch A

```

set protocols vrrp interface vlan-100 vif vlan-100 vrid 2 ip 172.168.30.1
set protocols vrrp interface vlan-100 vif vlan-100 vrid 2 load-balance disable false

```

Configure on Switch B

Configure Mlag on Switch B

```

set interface aggregate-ethernet ae1 aggregated-ether-options lACP enable true
set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id 2
set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer 192.168.30.1
peer-link ae3
set interface aggregate-ethernet ae1 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae1 family ethernet-switching vlan members 100
set interface aggregate-ethernet ae3 family ethernet-switching native-vlan-id 300
set interface aggregate-ethernet ae3 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae3 family ethernet-switching vlan members 100
set interface gigabit-ethernet qe-1/1/49 ether-options 802.3ad ae1
set interface gigabit-ethernet qe-1/1/50 family ethernet-switching native-vlan-id 500
set interface gigabit-ethernet qe-1/1/51 ether-options 802.3ad ae3
set interface gigabit-ethernet qe-1/1/53 ether-options 802.3ad ae3
set protocols static route 172.168.60.0/24 next-hop 192.168.50.2
set vlan-interface interface vlan-300 vif vlan-300 address 192.168.30.2 prefix-length 24
set vlan-interface interface vlan-100 vif vlan-100 address 172.168.30.3 prefix-length 24
set vlan-interface interface vlan-500 vif vlan-500 address 192.168.50.1 prefix-length 24
set vlans vlan-id 100 l3-interface vlan-100
set vlans vlan-id 300 l3-interface vlan-300
set vlans vlan-id 500 l3-interface vlan-500

```

Configure Vrrp on Switch B

```
set protocols vrrp interface vlan-100 vif vlan-100 vrid 2 ip 172.168.30.1
set protocols vrrp interface vlan-100 vif vlan-100 vrid 2 load-balance disable false
```

Configure on Switch C

Configure Mlag on Switch C

```
set interface aggregate-ethernet ae2 aggregated-ether-options lacp enable true
set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id 3
set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer 192.168.70.2
peer-link ae4
set interface aggregate-ethernet ae2 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae2 family ethernet-switching vlan members 200
set interface aggregate-ethernet ae4 family ethernet-switching native-vlan-id 700
set interface aggregate-ethernet ae4 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae4 family ethernet-switching vlan members 200
set interface gigabit-ethernet qe-1/1/49 ether-options 802.3ad ae2
set interface gigabit-ethernet qe-1/1/50 family ethernet-switching native-vlan-id 800
set interface gigabit-ethernet qe-1/1/51 ether-options 802.3ad ae4
set interface gigabit-ethernet qe-1/1/52 ether-options 802.3ad ae4
set protocols spanning-tree enable false
set protocols static route 172.168.30.0/24 next-hop 192.168.41.2
set vlan-interface interface vlan-200 vif vlan-200 address 172.168.60.2 prefix-length 24
set vlan-interface interface vlan-700 vif vlan-700 address 192.168.70.1 prefix-length 24
set vlan-interface interface vlan-800 vif vlan-800 address 192.168.41.1 prefix-length 24
set vlans vlan-id 200 l3-interface vlan-200
set vlans vlan-id 700 l3-interface vlan-700
set vlans vlan-id 800 l3-interface vlan-800
```

Configure Vrrp on Switch C

```
set protocols vrrp interface vlan-200 vif vlan-200 vrid 3 ip 172.168.60.1
set protocols vrrp interface vlan-200 vif vlan-200 vrid 3 load-balance disable false
```

Configure on Switch D

Configure Mlag on Switch D

```
set interface aggregate-ethernet ae2 aggregated-ether-options lacp enable true
set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id 3
set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer 192.168.70.1
peer-link ae4
set interface aggregate-ethernet ae2 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae2 family ethernet-switching vlan members 200
set interface aggregate-ethernet ae4 family ethernet-switching native-vlan-id 700
set interface aggregate-ethernet ae4 family ethernet-switching port-mode trunk
set interface aggregate-ethernet ae4 family ethernet-switching vlan members 200
set interface gigabit-ethernet qe-1/1/49 ether-options 802.3ad ae2
set interface gigabit-ethernet qe-1/1/50 family ethernet-switching native-vlan-id 900
set interface gigabit-ethernet qe-1/1/51 ether-options 802.3ad ae4
set interface gigabit-ethernet qe-1/1/52 ether-options 802.3ad ae4
```

```

set protocols spanning-tree enable false
set protocols static route 172.168.30.0/24 next-hop 192.168.51.2
set vlan-interface interface vlan-200 vif vlan-200 address 172.168.60.3 prefix-length 24
set vlan-interface interface vlan-700 vif vlan-700 address 192.168.70.2 prefix-length 24
set vlan-interface interface vlan-900 vif vlan-900 address 192.168.51.1 prefix-length 24
set vlans vlan-id 200 l3-interface vlan-200
set vlans vlan-id 700 l3-interface vlan-700
set vlans vlan-id 900 l3-interface vlan-900

```

Configure Vrrp on Switch D

```

set protocols vrrp interface vlan-200 vif vlan-200 vrid 3 ip 172.168.60.1
set protocols vrrp interface vlan-200 vif vlan-200 vrid 3 load-balance disable false

```

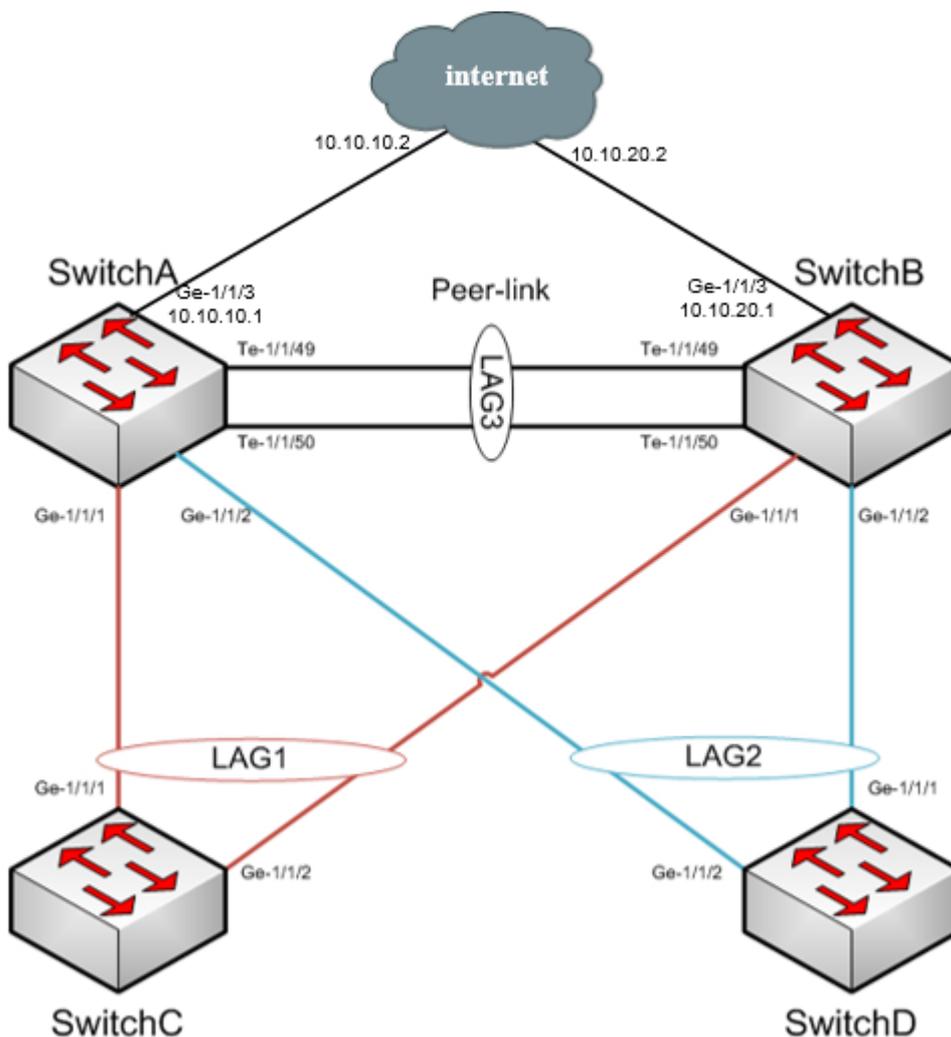
Configuring a MLAG with vxlan Example

Figure 3 illustrates MLAG configured between Switch A and Switch B, the MLAG connections between the neighboring switches, and two Network Devices.

The MLAG switches connect through a LACP LAG to Switch C.

The MLAG switches connect through a LACP LAG to Switch D.

Figure 3



- 01 Configuring mlag in Switch A and Switch B
 - Configuring Lag in Switch C and Switch D
- 02 Configuring Lag in Switch C and Switch D
- 03 Configuring vxlan in Switch A and Switch B

01 Configuring mlag in Switch A and Switch B

Switch A Configuration

```

admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring an Aggregation Interface to VLAN Members

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 13-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching
native-vlan-id 4094
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members

```

```
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 Interface IP Address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.1 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the Domain-id and System-id for the MLAG Domain

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id
1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the Peer IP Address and the Peer-link for the MLAG Domain Peer

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer
10.10.0.2 peer-link "ae3"
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer
10.10.0.2 peer-link "ae3"
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Switch B Configuration

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable
true
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options lacp enable
true
```

```
admin@XorPlus# set interface aggregate-ethernet ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 ether-options 802.3ad ae3
admin@XorPlus# set interface gigabit-ethernet te-1/1/50 ether-options 802.3ad ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring an Aggregation Interface to VLAN Members

```
admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# set vlans vlan-id 4094 l3-interface 4094
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
16
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching
native-vlan-id 4094
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ae3 family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the L3 Interface IP Address

```
admin@XorPlus# set vlan-interface interface 4094 vif 4094 address 10.10.0.2 prefix-length
24
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring the Domain-id and System-id for the MLAG Domain

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag domain-id  
1  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#  
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag domain-id  
2  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.
```

Configuring the Peer IP Address and the Peer-link for the MLAG Domain Peer

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options mlag peer  
10.10.0.1 peer-link "ae3"  
admin@XorPlus# set interface aggregate-ethernet ae2 aggregated-ether-options mlag peer  
10.10.0.1 peer-link "ae3"  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

- Configuring Lag in Switch C and Switch D

Configuring Lag in Switch C and Switch D

Switch C and Switch D Configuration

```
admin@XorPlus# set interface aggregate-ethernet ae1 aggregated-ether-options lacp enable  
true  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#  
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 ether-options 802.3ad ae1  
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 ether-options 802.3ad ae1  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# set vlans vlan-id 15
admin@XorPlus# set vlans vlan-id 16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface aggregate-ethernet ael family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ael family ethernet-switching vlan members
15
admin@XorPlus# set interface aggregate-ethernet ael family ethernet-switching vlan members
16
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the vxlan in Switch A and Switch B

Switch A

```

set vlan-interface loopback address 10.10.10.1 prefix-length 32
set vxlans source-interface loopback address 10.10.10.1
set vxlans vni 100010 interface ael vlan 15
set vxlans vni 100010 flood vtep 20.20.20.1
set protocols static route 20.20.20.1/24 next-hop 10.10.10.2

```

Switch B

```

set vlan-interface loopback address 10.10.10.1 prefix-length 32
set vxlans source-interface loopback address 10.10.10.1
set vxlans vni 100010 interface ael vlan 15
set vxlans vni 100010 flood vtep 20.20.20.1
set protocols static route 20.20.20.1/24 next-hop 10.10.20.2

```

Q-in-Q Basic Port Configuration

Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer2 Ethernet connection between two customer sites. Q-in-Q tunneling can also be used to segregate or bundle customer traffic into fewer VLANs, or different VLANs, by adding another layer of 802.1Q tags.

Q-in-Q tunneling is useful when there are overlapping VLAN IDs because the 802.1Q VLAN tags are prepended by the service VLAN tag. The L2/L3 implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

The Q-in-Q tunneling external mode belongs to basic Q-in-Q, while the Q-in-Q tunneling internal mode belongs to selective Q-in-Q.

Configuring the Q-in-Q Tunneling Internal/External Mode

By default, Q-in-Q is disabled. User can enable it as shown below:

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling mode
external
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Q-in-Q Tunneling to Map Ingress VLANs to Service VLANs

Selective Q-in-Q tunneling allows user to add different customer VLAN tags based on different service VLAN tags.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# set vlans dot1q-tunneling ingress t1 from untag enabled true
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then customer-vlan 10
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then service-vlan 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
ingress t1
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling ingress t2 from one-tag customer-vlan-list 20
admin@XorPlus# set vlans dot1q-tunneling ingress t2 then service-vlan 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling ingress t2
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling ingress t3 from one-tag customer-vlan-list 30
admin@XorPlus# set vlans dot1q-tunneling ingress t3 then service-vlan 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling ingress t3
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1 dot1q-tunneling
Dot1q Tunneling Mode: none, Ether Type: 0x8100
Ingress: t1
Untagged-type Enabled: true
One-tagged-type Customer Vlan:
Double-tagged-type Service Vlan: 0
New Service Vlan: 100
New Customer Vlan: 10
Ingress: t2
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 20
Double-tagged-type Service Vlan: 0

```

```

New Service Vlan: 200
New Customer Vlan: 0
Ingress: t3
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 30
Double-tagged-type Service Vlan: 0
New Service Vlan: 300
New Customer Vlan: 0
admin@XorPlus#

```

Configuring Q-in-Q Tunneling Egress Pop Service VLANs

Selective Q-in-Q tunneling allows user to delete different customer VLAN tags based on different service VLAN tags.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t1 from customer-vlan 10
admin@XorPlus# set vlans dot1q-tunneling egress t1 from service-vlan 100
admin@XorPlus# set vlans dot1q-tunneling egress t1 then action none
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t1
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t2 from customer-vlan 20
admin@XorPlus# set vlans dot1q-tunneling egress t2 from service-vlan 200
admin@XorPlus# set vlans dot1q-tunneling egress t2 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t2
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t3 from customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling egress t3 from service-vlan 300
admin@XorPlus# set vlans dot1q-tunneling egress t3 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t3
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
Egress: t1
Service Vlan: 100
Customer Vlan: 10
Action: Strip both tags

```

```

Egress: t2
Service Vlan: 200
Customer Vlan: 20
Action: Retain the customer vlan tag
Egress: t3
Service Vlan: 300
Customer Vlan: 30
Action: Retain the customer vlan tag
admin@XorPlus#

```

Q-in-Q Configuration Example

The configuration of Q-in-Q is shown in Fig. 4-2.

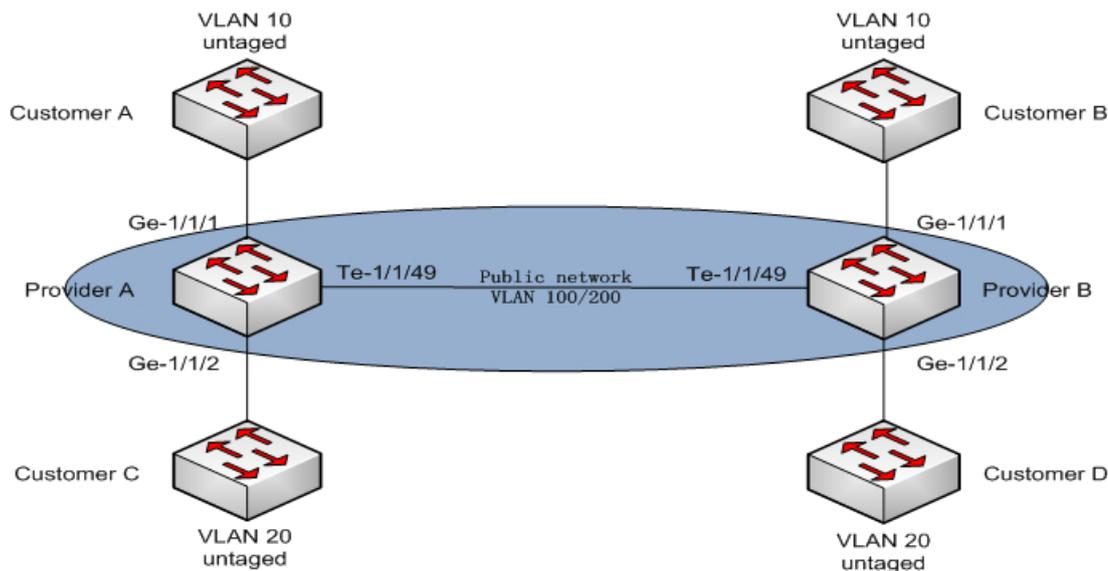


Figure 4-2. Q-in-Q configuration.

Configuration on Provider A

Configure VLAN 100 as the default VLAN of Gigabit Ethernet ge-1/1/1, and enable the Q-in-Q tunneling internal mode on Gigabit Ethernet ge-1/1/1.

Then, configure the untagged frames received by the port with the customer VLAN tag 30 and service VLAN tag 100.

Finally, configure the customer VLAN tag 10 frames received by the port with the service VLAN tag 100.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 100
admin@XorPlus# set vlans dot1q-tunneling ingress t1 from untag enabled true
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then service-vlan 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
ingress t1
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling ingress t2 from one-tag customer-vlan-list 10
admin@XorPlus# set vlans dot1q-tunneling ingress t2 then service-vlan 100

```

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling ingress t2
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t3 from customer-vlan 10
admin@XorPlus# set vlans dot1q-tunneling egress t3 from service-vlan 100
admin@XorPlus# set vlans dot1q-tunneling egress t3 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t3
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t4 from customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling egress t4 from service-vlan 100
admin@XorPlus# set vlans dot1q-tunneling egress t4 then action none
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
Ingress: t1
Untagged-type Enabled: true
One-tagged-type Customer Vlan:
Double-tagged-type Service Vlan: 0
New Service Vlan: 100
New Customer Vlan: 30
Ingress: t2
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 10
Double-tagged-type Service Vlan: 0
New Service Vlan: 100
New Customer Vlan: 0
Egress: t3
Service Vlan: 100
Customer Vlan: 10
Action: Retain the customer vlan tag
Egress: t4
Service Vlan: 100
Customer Vlan: 30
Action: Strip both tags
admin@XorPlus#

```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet ge-1/1/2, and enable the Q-in-Q tunneling internal mode on Gigabit Ethernet ge-1/1/2.

Then, configure the untagged frames received by the port with the customer VLAN tag 30 and service VLAN tag 200.

Finally, configure the customer VLAN tag 20 frames received by the port with the service VLAN Tag 200.

```

admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 200
admin@XorPlus# set vlans dot1q-tunneling ingress t5 from untag enabled true
admin@XorPlus# set vlans dot1q-tunneling ingress t5 then customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling ingress t5 then service-vlan 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
ingress t5
admin@XorPlus# commit

```

```

admin@XorPlus# set vlans dot1q-tunneling ingress t6 from one-tag customer-vlan-list 20
admin@XorPlus# set vlans dot1q-tunneling ingress t6 then service-vlan 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling ingress t6
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t7 from customer-vlan 20
admin@XorPlus# set vlans dot1q-tunneling egress t7 from service-vlan 200
admin@XorPlus# set vlans dot1q-tunneling egress t7 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
egress t7
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t8 from customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling egress t8 from service-vlan 200
admin@XorPlus# set vlans dot1q-tunneling egress t8 then action none
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
egress t8
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/2 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
Ingress: t5
Untagged-type Enabled: true
One-tagged-type Customer Vlan:
Double-tagged-type Service Vlan: 0
New Service Vlan: 200
New Customer Vlan: 30
Ingress: t6
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 20
Double-tagged-type Service Vlan: 0
New Service Vlan: 200
New Customer Vlan: 0
Egress: t7
Service Vlan: 200
Customer Vlan: 20
Action: Retain the customer vlan tag
Egress: t8
Service Vlan: 200
Customer Vlan: 30
Action: Strip both tags
admin@XorPlus#

```

Configure VLAN 100/200 as the trunk port of Gigabit Ethernet te-1/1/49, and enable the Q-in-Q tunneling internal mode.

```

admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
dot1q-tunneling modeinternal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet te-1/1/49 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100

```

Configuration on Provider B

Configure VLAN 100 as the default VLAN of Gigabit Ethernet ge-1/1/1, and enable the Q-in-Q tunneling internal mode on Gigabit Ethernet ge-1/1/1.

Then, configure the untagged frames received by the port with the customer VLAN tag 30 and service VLAN tag 100.

Finally, configure the customer VLAN tag 10 frames received by the port with the service VLAN tag 100.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 100
admin@XorPlus# set vlans dot1q-tunneling ingress t1 from untag enabled true
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling ingress t1 then service-vlan 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
ingress t1
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling ingress t2 from one-tag customer-vlan-list 10
admin@XorPlus# set vlans dot1q-tunneling ingress t2 then service-vlan 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling ingress t2
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t3 from customer-vlan 10
admin@XorPlus# set vlans dot1q-tunneling egress t3 from service-vlan 100
admin@XorPlus# set vlans dot1q-tunneling egress t3 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t3
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t4 from customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling egress t4 from service-vlan 100
admin@XorPlus# set vlans dot1q-tunneling egress t4 then action none
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling
egress t4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
Ingress: t1
Untagged-type Enabled: true
One-tagged-type Customer Vlan:
Double-tagged-type Service Vlan: 0
New Service Vlan: 100
New Customer Vlan: 30
Ingress: t2
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 10
Double-tagged-type Service Vlan: 0
New Service Vlan: 100
New Customer Vlan: 0
Egress: t3
Service Vlan: 100
Customer Vlan: 10
Action: Retain the customer vlan tag
Egress: t4
Service Vlan: 100

```

```
Customer Vlan: 30
Action: Strip both tags
admin@XorPlus#
```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet ge-1/1/2, and enable the Q-in-Q tunneling internal mode on Gigabit Ethernet 1/1/2.

Then, configure the untagged frames received by the port with the customer VLAN tag 30 and service VLAN tag 200.

Finally, configure the customer VLAN tag 20 frames received by the port with the service VLAN Tag 200.

```
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 200
admin@XorPlus# set vlans dot1q-tunneling ingress t5 from untag enabled true
admin@XorPlus# set vlans dot1q-tunneling ingress t5 then customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling ingress t5 then service-vlan 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
ingress t5
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling ingress t6 from one-tag customer-vlan-list 20
admin@XorPlus# set vlans dot1q-tunneling ingress t6 then service-vlan 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling ingress t6
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t7 from customer-vlan 20
admin@XorPlus# set vlans dot1q-tunneling egress t7 from service-vlan 200
admin@XorPlus# set vlans dot1q-tunneling egress t7 then action one
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
egress t7
admin@XorPlus# commit
admin@XorPlus# set vlans dot1q-tunneling egress t8 from customer-vlan 30
admin@XorPlus# set vlans dot1q-tunneling egress t8 from service-vlan 200
admin@XorPlus# set vlans dot1q-tunneling egress t8 then action none
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling
egress t8
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
dot1q-tunneling mode
internal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/2 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
Ingress: t5
Untagged-type Enabled: true
One-tagged-type Customer Vlan:
Double-tagged-type Service Vlan: 0
New Service Vlan: 200
New Customer Vlan: 30
Ingress: t6
Untagged-type Enabled: false
One-tagged-type Customer Vlan: 20
Double-tagged-type Service Vlan: 0
New Service Vlan: 200
New Customer Vlan: 0
Egress: t7
Service Vlan: 200
Customer Vlan: 20
Action: Retain the customer vlan tag
```

```
Egress: t8
Service Vlan: 200
Customer Vlan: 30
Action: Strip both tags
admin@XorPlus#
```

Configure VLAN 100/200 as the trunk port of Gigabit Ethernet te-1/1/49, and enable the Q-in-Q tunneling internal mode.

```
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
dot1q-tunneling modeinternal
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet te-1/1/49 dot1q-tunneling
Dot1q Tunneling Mode: internal, Ether Type: 0x8100
admin@XorPlus#
```

MSTP Configuration

802.1D, 802.1w, and 802.1s are spanning tree protocols that can avoid the loop in Layer2. You can configure the parameters of MSTP, including bridge-priority, forward-delay, max-age, and hello-time interval.

Enabling Spanning Tree Mode in MSTP

```
admin@XorPlus# set protocols spanning-tree force-version 3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Basic Global Parameters of MSTP

When configuring global parameters, make sure to set the forward delay to greater than $\text{Max-Age}/2 + 1$, or the commit will fail.

```
admin@XorPlus# set protocols spanning-tree mstp bridge-priority 4096
admin@XorPlus# set protocols spanning-tree mstp forward-delay 20
admin@XorPlus# set protocols spanning-tree mstp hello-time 2
admin@XorPlus# set protocols spanning-tree mstp max-age 20
admin@XorPlus# set protocols spanning-tree mstp max-hops 8
admin@XorPlus# set protocols spanning-tree mstp configuration-name test1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show spanning-tree mstp bridge
```

```
Bridge Spanning Tree Parameters
Enabled Protocol: MSTP
Root ID: 4096.08:9e:01:39:1a:fe
External Root Path Cost: 0
CIST Regional Root ID: 4096.08:9e:01:39:1a:fe
Root Port:
CIST Internal Root Path Cost: 0
Hello Time: 2
Maximum Age: 20
Forward Delay: 20
Remaining Hops: 8
Bridge Configuration Name: test1
Bridge Configuration Digest: ac36177f50283cd4b83821d8ab26de62
Number of Topology Changes: 13
Time Since Last Topology Change: 0 days 00:00:31
Local Parameters
Bridge ID: 4096.08:9e:01:39:1a:fe
Hello Time: 2
Maximum Age: 20
Forward Delay: 20
Remaining Hops: 8
admin@XorPlus#
admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1
admin@XorPlus# set protocols spanning-tree mstp msti 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 300
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 400
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree mstp bridge
Bridge Spanning Tree Parameters
Enabled Protocol: MSTP
Root ID: 4096.08:9e:01:39:1a:fe
External Root Path Cost: 0
CIST Regional Root ID: 4096.08:9e:01:39:1a:fe
Root Port:
CIST Internal Root Path Cost: 0
Hello Time: 2
Maximum Age: 20
Forward Delay: 20
Remaining Hops: 8
Bridge Configuration Name: test1
Bridge Configuration Digest: 8b5d98ca042bad0d7fa5f18744f4755d
Msti 1 Member VLANs:
100, 200,
Msti 2 Member VLANs:
300, 400,
Number of Topology Changes: 14
Time Since Last Topology Change: 0 days 00:02:49
Local Parameters
Bridge ID: 4096.08:9e:01:39:1a:fe
Hello Time: 2
Maximum Age: 20
Forward Delay: 20
```

```
Remaining Hops: 8
admin@XorPlus#
```

Configuring MSTP Interface Parameters

```
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 external-path-cost 30000
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 internal-path-cost 10000
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 edge true
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 mode point-to-point
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 port-priority 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show spanning-tree mstp interface
Spanning Tree Interface Parameters for Instance 0
Interface Port ID Designated Designated Bridge Ext Port Int Port State Role
Port ID ID Cost Cost
-----
-----
ge-1/1/1 96.1 96.1 8192.08:9e:01:39:1a:fe 30000 10000 FORWARDING EDGE
```

Configuring the BPDU Filter

The BPDU filter prevents the bridge from using BPDUs for STP calculations. The switch then ignores any BPDUs that it receives.

```
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 bpdu-filter true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BPDU Root Guard

If a switch port receives a higher bridge-priority BPDU, it will ignore the BPDU and keep the current root-bridge as the root-bridge.

```
admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 root-guard true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BPDU TCN-Guard

When a port is configured with TCN-guard, the port does not process or propagate any topology change information received on the configured port.

```

admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 tcn-guard true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring MSTP BPDU-Guard

When a port is configured with mstp bpdu-guard, the port which should not receive bpdu receives bpdu messages, it will be set to err-discard True and down.

```

admin@XorPlus# set protocols spanning-tree mstp interface ge-1/1/1 bpdu-guard true
admin@XorPlus# co
Commit OK.
Save done.

```

If the port is down by mstp bpdu-guard, the port status as follows.

```

admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1
Physical interface: ge-1/1/1, Enabled, error-discard True, Physical link is Down
Interface index: 1, Mac Learning Enabled
Description: User Port
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Storm control ratio Broadcast: 1%
Interface rate limit ingress:0, egress:0
Current address: c4:39:3a:ff:2d:c1, Hardware address: c4:39:3a:ff:2d:c1
Traffic statistics:
5 sec input rate 0 bits/sec, 0 packets/sec
5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....132
Output Packets.....4733
Input Octets.....11266
Output Octets.....1664371

```

When the port is down by mstp bpdu-guard, the port will be up after you delete bpdu-guard and disable, enable the port.

```

admin@XorPlus# delete protocols spanning-tree mstp interface ge-1/1/1 bpdu-guard
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable true
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1
Physical interface: ge-1/1/1, Enabled, error-discard False, Physical link is Up
Interface index: 1, Mac Learning Enabled
Description: User Port
Link-level type: Ethernet, MTU: 1514, Speed: 1Gb/s, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Disabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Storm control ratio Broadcast: 1%

```

```

Interface rate limit  ingress:0, egress:0
Current address: c4:39:3a:ff:2d:c1, Hardware address: c4:39:3a:ff:2d:c1
Traffic statistics:
 5 sec input rate 0 bits/sec, 0 packets/sec
 5 sec output rate 688 bits/sec, 0 packets/sec
Input Packets.....132
Output Packets.....4736
Input Octets.....11266
Output Octets.....1664755

```

Disabling/Enabling MSTP

If you disable MSTP, the port will stay in forwarding status and cease to send BPDUs.

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree mstp interface
Spanning Tree Interface Parameters for Instance 0
Interface Port ID Designated Designated Bridge Ext Port Int Port State Role
Port ID ID Cost Cost
-----
-----
ge-1/1/1 96.1 96.1 8192.08:9e:01:39:1a:fe 30000 10000 FORWARDING MSTP DISABLED
ge-1/1/2 128.2 128.2 8192.08:9e:01:39:1a:fe 20000 20000 FORWARDING MSTP DISABLED
ge-1/1/13 128.13 128.13 8192.08:9e:01:39:1a:fe 20000 20000 FORWARDING MSTP DISABLED
admin@XorPlus# set protocols spanning-tree enable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree mstp interface
Spanning Tree Interface Parameters for Instance 0
Interface Port ID Designated Designated Bridge Ext Port Int Port State Role
Port ID ID Cost Cost
-----
-----
ge-1/1/1 96.1 96.1 8192.08:9e:01:39:1a:fe 30000 10000 FORWARDING EDGE
ge-1/1/2 128.2 128.2 8192.08:9e:01:39:1a:fe 20000 20000 FORWARDING EDGE
ge-1/1/13 128.13 128.13 8192.08:9e:01:39:1a:fe 20000 20000 FORWARDING DESIGNATED

```

MSTP Configuration Example

There are two examples of MSTP configuration. In our first example, VLAN 100 is mapped to MSTI-1 and VLAN 200 is mapped to MSTI-2. The entire topology belongs to only one MSTP domain named **region1**. Switch A is the root of the network.

To achieve load balancing, VLAN 100 should be in MSTI-1 (Fig. 4-4), and VLAN 200 should be in MSTI-2 (Fig. 4-5).

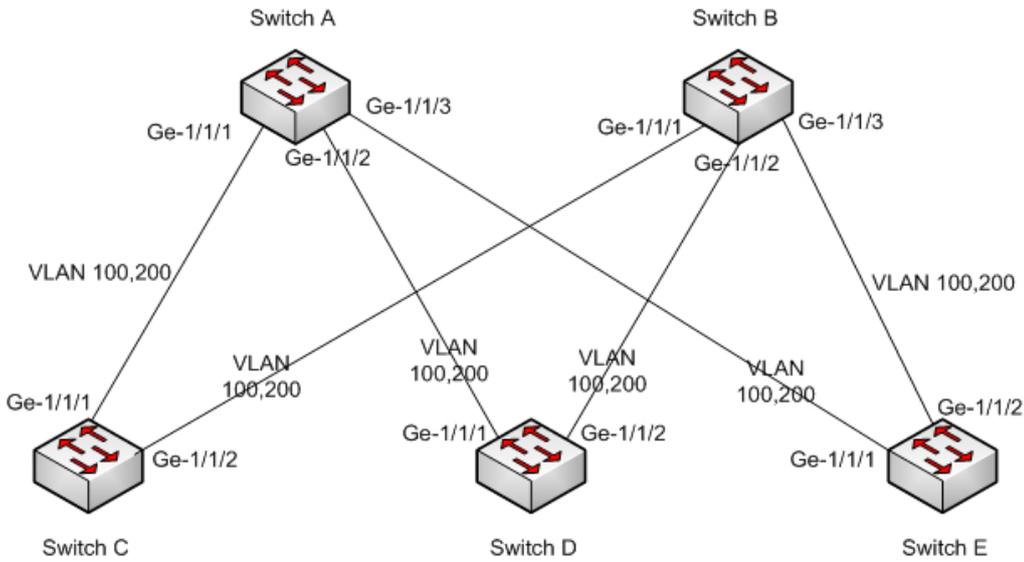


Figure 4-3. MSTP configuration.

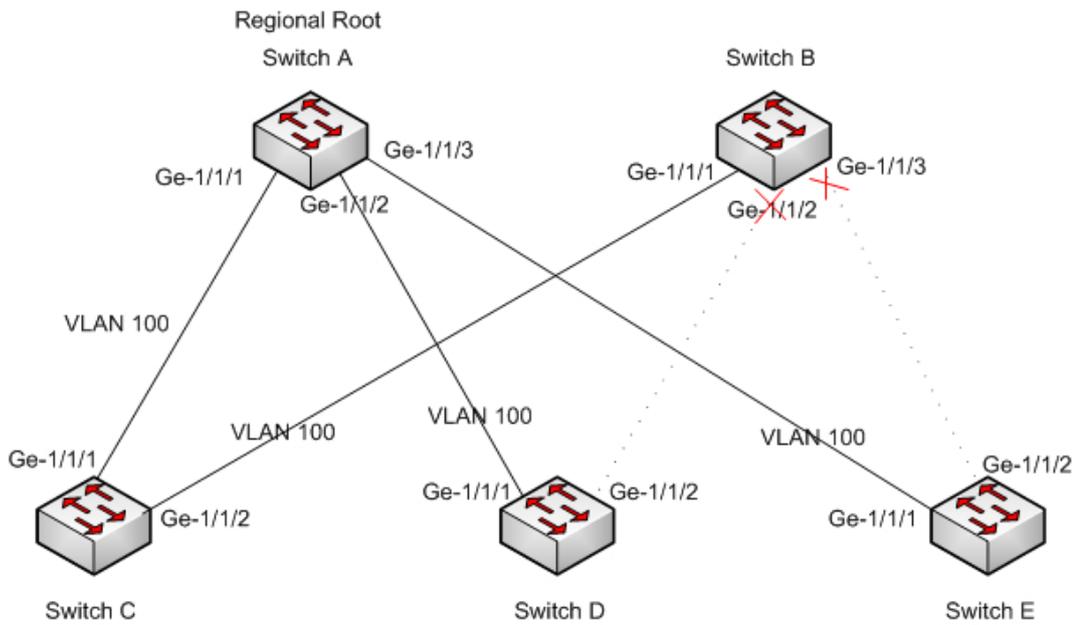


Figure 4-4. MSTI-1 topology for VLAN 100.

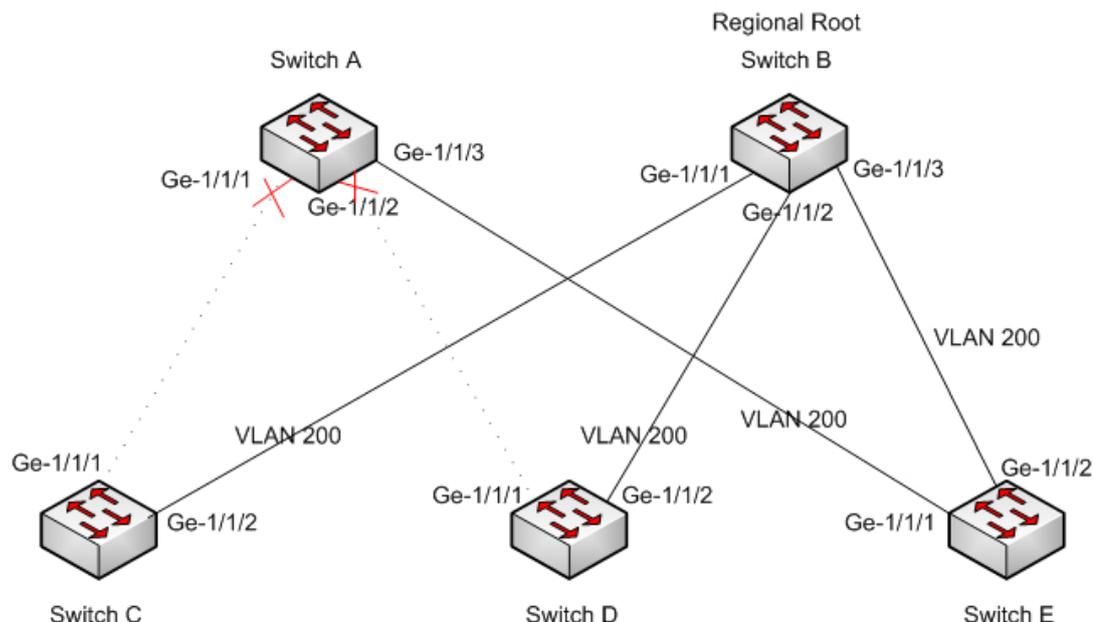


Figure 4-5. MSTI-2 topology for VLAN 200.

Configuring Switch A, Example 1

For Switch A, configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100 and VLAN 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
To make sure that Switch A is the root of the network and the regional root of MSTI-1,
configure it as the higher priority.

```

```

admin@XorPlus# set protocols spanning-tree mstp bridge-priority 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1 bridge-priority 4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B, Example 1

Configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100 and VLAN 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch B is the regional root of MSTI-2 and that ge-1/1/2 and ge-1/1/3 are in blocking status in MSTI-1, configure a higher MSTI-2 priority and a large value for internal-path-cost in MSTI-1.

```

admin@XorPlus# set protocols spanning-tree mstp msti 2 bridge-priority 4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/2 cost 10000000
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/3 cost 10000000
admin@XorPlus# commit

```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C, Example 1

Configure ge-1/1/1~ge-1/1/2 as trunk ports and as members of VLAN 100 and VLAN 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To set ge-1/1/1 and ge-1/1/2 in forwarding status in MSTI-1, configure a lower value for internal-path-cost.

To set ge-1/1/1 in blocking status in MSTI-2, configure a higher value for internal-path-cost.

```

admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/1 cost 1000
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/2 cost 1000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 2 interface ge-1/1/1 cost 100000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch D, Example 1

Configure ge-1/1/1~ge-1/1/2 as trunk ports and as members of VLAN 100 and VLAN 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To set ge-1/1/1 in blocking status in MSTI-2 and ge-1/1/2 in blocking status in MSTI-1, configure a large value for internal-path-cost.

```

admin@XorPlus# set protocols spanning-tree mstp msti 2 interface ge-1/1/1 cost 10000000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/2 cost 10000000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch E, Example 1

Configure ge-1/1/1~ge-1/1/2 as trunk ports and as members of VLAN 100 and VLAN 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200

```

```

admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To set ge-1/1/1 and ge-1/1/2 in forwarding status in MSTI-2, configure a lower value for internal-path-cost.

To set ge-1/1/2 in blocking status in MSTI-1, configure a large value for internal-path-cost.

```

admin@XorPlus# set protocols spanning-tree mstp msti 2 interface ge-1/1/1 cost 1000
admin@XorPlus# set protocols spanning-tree mstp msti 2 interface ge-1/1/2 cost 1000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/2 cost 10000000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

In the second example, there are two regions. In region 1, VLAN 100 is mapped to MSTI-1, VLAN 200 is mapped to MSTI-2, and VLAN 300 is mapped to MSTI-3. In region 2, VLAN 200 is mapped to MSTI-2 and VLAN 400 is mapped to MSTI-4. Switch A is the root of the entire network. The topologies of the VLANs are presented in Fig. 4-6 through 4-10.

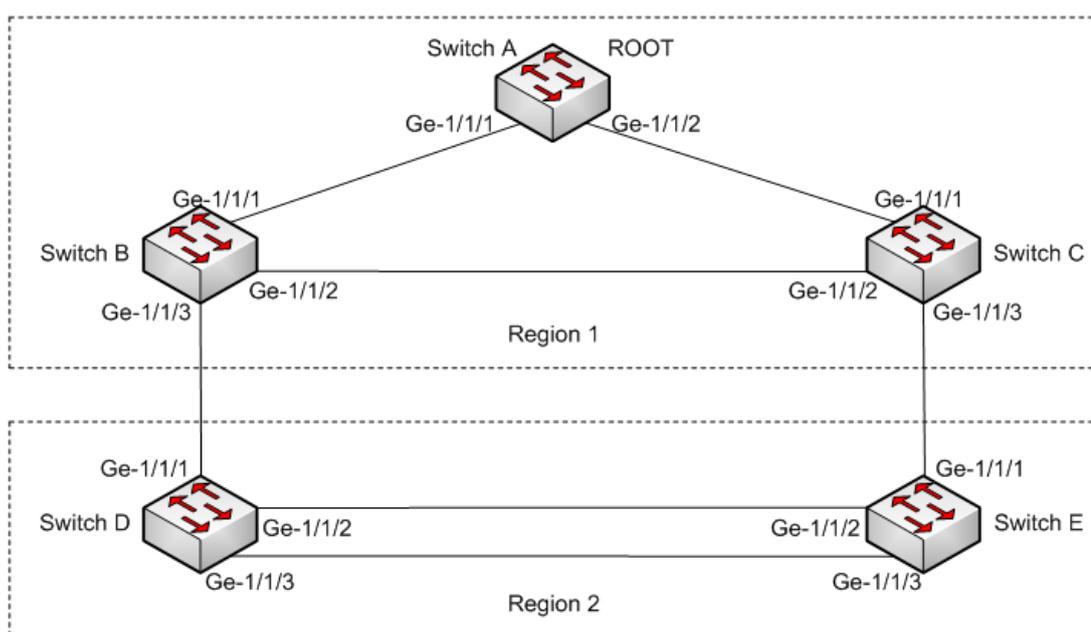


Figure 4-6. MSTP configuration.

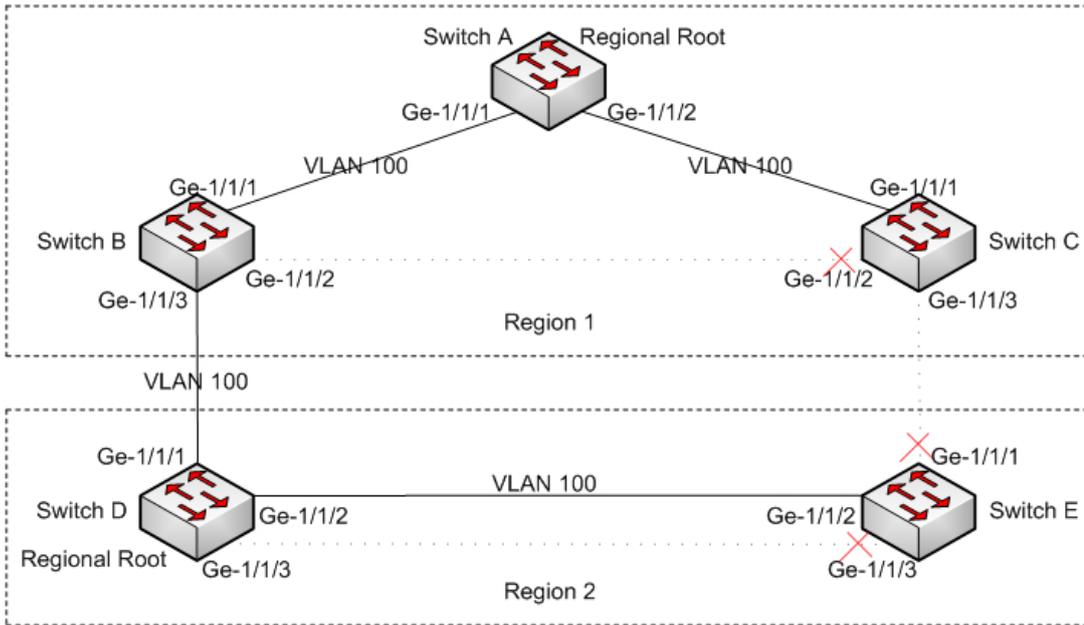


Figure 4-7. Topology for VLAN 100.

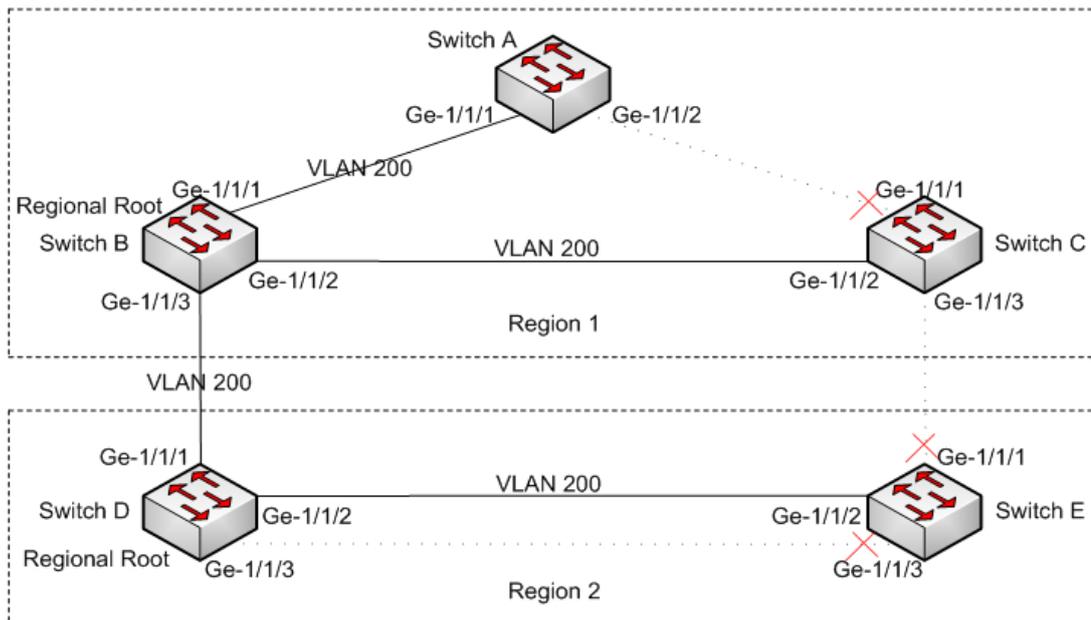


Figure 4-8. Topology for VLAN 200.

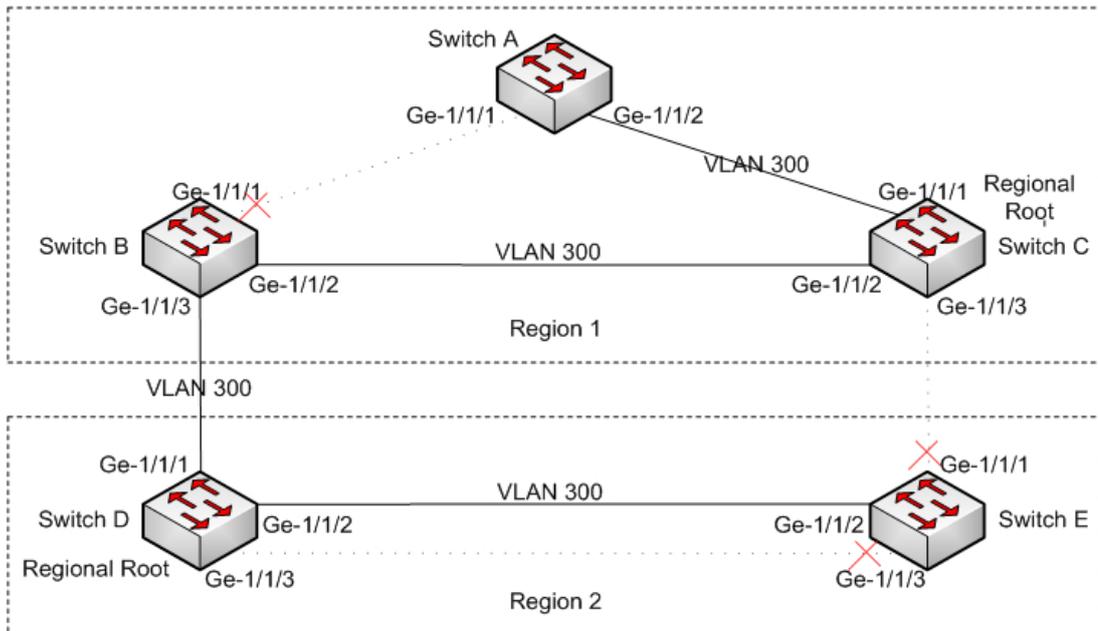


Figure 4-9. Topology for VLAN 300.

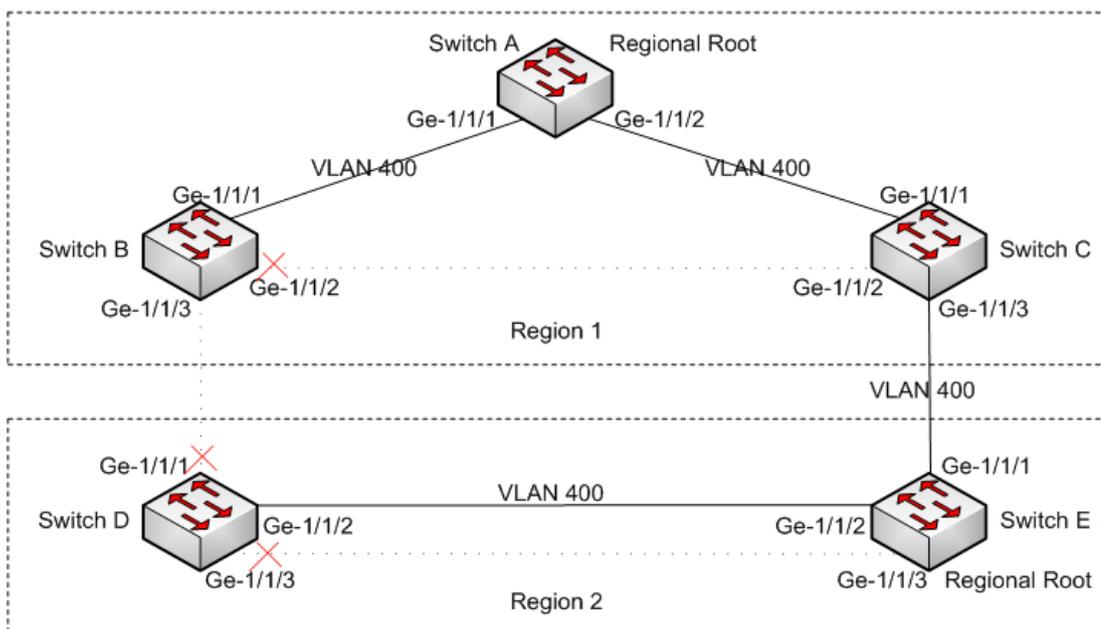


Figure 4-10. Topology for VLAN 400.

Configuring Switch A, Example 2

For Switch A, configure ge-1/1/1~ge-1/1/2 as trunk ports and as members of VLAN 100, VLAN 200, VLAN 300, and VLAN 400.

```
admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
```

```

trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 400
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 3 vlan 300
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
To verify that Switch A is the root of the network and the regional root of MSTI-1,
configure it as the higher priority.
admin@XorPlus# set protocols spanning-tree mstp bridge-priority 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 1 bridge-priority 4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B, Example 2

Configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100, VLAN 200, VLAN 300, and VLAN 400.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan

```

```

members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 400
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 3 vlan 300
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch B is the regional root of MSTI-2 and that ge-1/1/1 is in blocking status in MSTI-3, configure a higher MSTI-2 priority and a large value for internal-path-cost in MSTI-3.

```

admin@XorPlus# set protocols mstp msti 2 bridge-priority 4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols mstp msti 3 interface ge-1/1/1 cost 10000000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C, Example 2

Configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100, VLAN 200, VLAN 300, and VLAN 400.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400

```

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 400
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 1 vlan 100
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 3 vlan 300
admin@XorPlus# set protocols spanning-tree mstp configuration-name region1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch C is the regional root of MSTI-3, ge-1/1/1 is in blocking status in MSTI-2, and ge-1/1/2 is in blocking status in MSTI-1, user should configure a higher MSTI-3 priority and large values for internal-path-costs of ge-1/1/1 in MSTI-2 and ge-1/1/2 in MSTI-1.

```

admin@XorPlus# set protocols spanning-tree mstp msti 3 bridge-priority 4096
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 2 interface ge-1/1/1 cost 10000000
admin@XorPlus# set protocols spanning-tree mstp msti 1 interface ge-1/1/2 cost 10000000
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch D, Example 2

Configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100, VLAN 200, VLAN 300, and VLAN 400.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 400
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 4 vlan 400
admin@XorPlus# set protocols spanning-tree mstp configuration-name region2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch D is the regional root of MSTI-2 and the root of CIST. Configure a higher MSTI-2 priority and bridge priority.

```

admin@XorPlus# set protocols spanning-tree mstp bridge-priority 16384
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree mstp msti 2 bridge-priority 4096
admin@XorPlus# commit

```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch E, Example 2

Configure ge-1/1/1~ge-1/1/3 as trunk ports and as members of VLAN 100, VLAN 200, VLAN 300, and VLAN 400.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 400
admin@XorPlus#
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree mstp msti 2 vlan 200
admin@XorPlus# set protocols spanning-tree mstp msti 4 vlan 400
admin@XorPlus# set protocols spanning-tree mstp configuration-name region2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch E is the regional root of MSTI-4. Configure a higher MSTI-4 priority.

```

admin@XorPlus# set protocols spanning-tree mstp msti 4 bridge-priority 4096
admin@XorPlus# commit

```

```
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

PVST Configuration

802.1D, 802.1w, and 802.1s are spanning tree protocols that avoid the loop in Layer 2. You can configure the parameters of PVST, including bridge-priority, forward-delay, max-age, and hello-time interval.

Enabling Spanning Tree Mode in PVST

```
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Basic VLAN Parameters of PVST

When configuring basic VLAN parameters, set the forward delay to greater than $\text{Max-Age}/2 + 1$, or the commit will fail.

```
admin@XorPlus# set protocols spanning-tree pvst vlan 2 bridge-priority 4096
admin@XorPlus# set protocols spanning-tree pvst vlan 2 forward-delay 20
admin@XorPlus# set protocols spanning-tree pvst vlan 2 hello-time 4
admin@XorPlus# set protocols spanning-tree pvst vlan 2 max-age 30
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree pvst bridge vlan 2
PVST Bridge Parameters for VLAN 2
Root Bridge: 4098.08:9e:01:61:65:71
Root Cost: 0
Root Port:
Hello Time: 4
Max Age: 30
Forward Delay: 20
Time Since Last Topology Change: 0 days 00:02:55
Local Parameters
Bridge ID: 4098.08:9e:01:61:65:71
Hello Time: 4
Maximum Age: 30
Forward Delay: 20
```

Configuring PVST Interface Parameters

```
admin@XorPlus# set protocols spanning-tree pvst vlan 2 interface ge-1/1/1 path-cost 555555
admin@XorPlus# set protocols spanning-tree pvst vlan 2 interface ge-1/1/1 port-priority 200
admin@XorPlus# commit
```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show spanning-tree pvst interface vlan 2
Rapid PVST+ Spanning Tree Interface Status for VLAN 2
Interface Port ID Designated Designated Bridge Port Cost State Role
Port ID ID
-----
ge-1/1/1 192.1 192.1 4098.08:9e:01:61:65:71 555555 FORWARDING EDGE

```

Configuring the Interface Mode

You can configure the interface mode as point-to-point or shared.

```

admin@XorPlus# set protocols spanning-tree pvst interface ge-1/1/1 mode point-to-point
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree pvst interface ge-1/1/1 mode shared
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring PVST BPDU-guard

When a port is configured with `pvst bpdu-guard`, the port which should not receive bpdu receives bpdu messages, it will be set to `err-discard True` and down.

```

admin@XorPlus# set protocols spanning-tree pvst interface ge-1/1/1 bpdu-guard true
admin@XorPlus# co
Commit OK.
Save done.

```

If the port is down by `pvst bpdu-guard`, the port status as follows.

```

admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1
Physical interface: ge-1/1/1, Enabled, error-discard True, Physical link is Down
Interface index: 1, Mac Learning Enabled
Description: User Port
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Enabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Storm control ratio Broadcast: 1%
Interface rate limit ingress:0, egress:0
Current address: c4:39:3a:ff:2d:c1, Hardware address: c4:39:3a:ff:2d:c1
Traffic statistics:
5 sec input rate 0 bits/sec, 0 packets/sec
5 sec output rate 0 bits/sec, 0 packets/sec
Input Packets.....132
Output Packets.....4733
Input Octets.....11266
Output Octets.....1664371

```

When the port is down by pvst bpdu-guard, the port will be up after you delete bpdu-guard and disable, enable the port.

```

admin@XorPlus# delete protocols spanning-tree pvst interface ge-1/1/1 bpdu-guard
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable true
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show interface gigabit-ethernet ge-1/1/1
Physical interface: ge-1/1/1, Enabled, error-discard False, Physical link is Up
Interface index: 1, Mac Learning Enabled
Description: User Port
Link-level type: Ethernet, MTU: 1514, Speed: 1Gb/s, Duplex: Full
Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Disabled
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Storm control ratio Broadcast: 1%
Interface rate limit ingress:0, egress:0
Current address: c4:39:3a:ff:2d:c1, Hardware address: c4:39:3a:ff:2d:c1
Traffic statistics:
 5 sec input rate 0 bits/sec, 0 packets/sec
 5 sec output rate 688 bits/sec, 0 packets/sec
Input Packets.....132
Output Packets.....4736
Input Octets.....11266
Output Octets.....1664755

```

Disabling/Enabling PVST on One VLAN

You can disable or enable the spanning tree protocol PVST on a single designated VLAN.

```

admin@XorPlus# set protocols spanning-tree pvst vlan 2 enable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree pvst bridge vlan 2
PVST Bridge Parameters for VLAN 2
Root Bridge: 32769.08:9e:01:61:65:71
Root Cost: 0
Root Port:
Hello Time: 2
Max Age: 20
Forward Delay: 15
Time Since Last Topology Change: 15804 days 23:00:11
Local Parameters
Bridge ID: 32769.08:9e:01:61:65:71
Hello Time: 2
Maximum Age: 20
Forward Delay: 15
admin@XorPlus# set protocols spanning-tree pvst vlan 2 enable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

```

admin@XorPlus#
admin@XorPlus# run show spanning-tree pvst bridge vlan 2
PVST Bridge Parameters for VLAN 2
Root Bridge: 4098.08:9e:01:61:65:71
Root Cost: 0
Root Port:
Hello Time: 4
Max Age: 30
Forward Delay: 20
Time Since Last Topology Change: 0 days 00:00:21
Local Parameters
Bridge ID: 4098.08:9e:01:61:65:71
Hello Time: 4
Maximum Age: 30
Forward Delay: 20
admin@XorPlus#

```

Disabling/Enabling PVST

You cannot disable the spanning tree protocol PVST with just the **enable false** command. To disable PVST, *first* configure the spanning tree mode in MSTP/RSTP/STP. Then, disable the spanning tree. After the spanning tree is disabled, the port will stay in "forwarding" status and cease to send BPDUs.

```

admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit Failed
102 Command failed Cannot disable spanning tree under PVST mode[
admin@XorPlus#
admin@XorPlus# exit discard
admin@XorPlus> configure
Entering configuration mode.
There are no other users in configuration mode.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree enable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols spanning-tree enable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show spanning-tree
Bridge Spanning Tree Parameters
Enabled Protocol: PVST
Root ID: 32769.08:9e:01:61:65:71
Root Path Cost: 0
Designated Bridge ID: 32769.08:9e:01:61:65:71
Root Port:
Hello Time: 2

```

```

Maximum Age: 20
Forward Delay: 15
Number of Topology Changes: 1
Time Since Last Topology Change: 0 days 00:00:09
Local Parameters
Bridge ID: 32769.08:9e:01:61:65:71
Hello Time: 2
Maximum Age: 20
Forward Delay: 15

```

PVST Configuration Example

The following topology is an example of a PVST configuration. Switches A and B are in the aggregation layer, and switches C and D are in the access layer. Configure switch A as the root bridge of VLAN 100 and VLAN 200, switch B as the root bridge of VLAN 300, and switch C as the root bridge of VLAN 400.

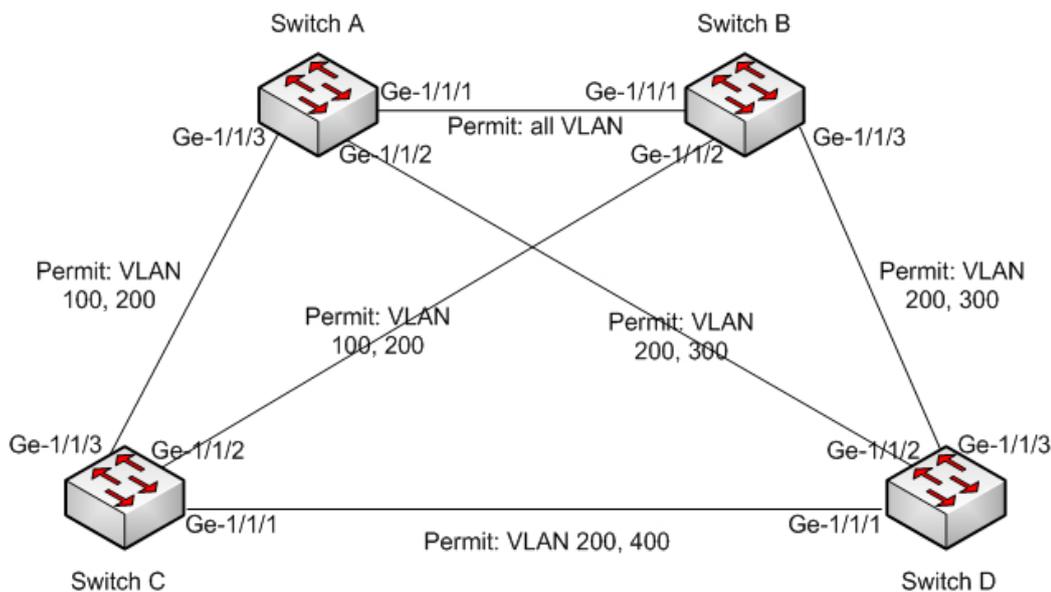


Figure 4-11. PVST configuration.

Configuring Switch A

For Switch A, configure ge-1/1/1~ge-1/1/3 as trunk ports; ge-1/1/1 as a member of VLANs 100, 200, 300, and 400; ge-1/1/2 as a member of VLANs 200 and 300; and ge-1/1/3 as a member of VLANs 100 and 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 400

```

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch A is the root bridge of VLANs 100 and 200, configure VLANs 100 and 200 as the higher priority.

```

admin@XorPlus# set protocols spanning-tree pvst vlan 100 bridge-priority 0
admin@XorPlus# set protocols spanning-tree pvst vlan 200 bridge-priority 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B

Configure ge-1/1/1~ge-1/1/3 as trunk ports, and ge-1/1/1 as a member of VLANs 100, 200, 300, and 400; ge-1/1/2 as a member of VLANs 100 and 200; and ge-1/1/3 as a member of VLANs 200 and 300.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk

```

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 300
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch B is the root bridge of VLAN 300, configure VLAN 300 as the higher priority.

```

admin@XorPlus# set protocols spanning-tree pvst vlan 300 bridge-priority 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure ge-1/1/1~ge-1/1/3 as trunk ports; ge-1/1/1 as a member of VLANs 200 and 400; ge-1/1/2 as a member of VLANs 100 and 200; and ge-1/1/3 as a member of VLANs 100 and 200.

```

admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan
members 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

To verify that Switch C is the root bridge of VLAN 400, configure VLAN 400 as the higher priority.

```
admin@XorPlus# set protocols spanning-tree pvst vlan 400 bridge-priority 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch D

Configure ge-1/1/1~ge-1/1/3 as trunk ports; ge-1/1/1 as a member of VLANs 200 and 400; ge-1/1/2 as a member of VLANs 200 and 300; and ge-1/1/3 as a member of VLANs 200 and 300.

```
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# set vlans vlan-id 300
admin@XorPlus# set vlans vlan-id 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan members 300
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan members 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching vlan members 300
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols spanning-tree force-version 4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Mirroring

You can configure one (1) mirror to analyze traffic. Configure the source/destination port (input/output port).

Note: The output (mirroring) port can belong to any VLAN and will not participate in Layer2 or Layer3 forwarding.

Configuring Mirroring to Analyze Traffic

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input egress ge-1/1/1
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input ingress ge-1/1/1

admin@XorPlus# set interface ethernet-switching-options analyzer 111 input egress ge-1/1/2
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input ingress ge-1/1/2
admin@XorPlus# set interface ethernet-switching-options analyzer 111 output ge-1/1/3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show analyzer 111
Analyzer name: 111
Output interface: <ge-1/1/3>
Ingress monitored interfaces: <ge-1/1/1><ge-1/1/2>
Egress monitored interfaces: <ge-1/1/1><ge-1/1/2>
admin@XorPlus#
```

Configuring Mirroring Guide

User can configure a port as a mirroring port, which can analyze the traffic of egress port or ingress port.

Configuring a Port as a Mirroring Port

When configuring mirroring, user must first configure the port as a mirror.

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 output ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

Configuring Mirroring on Egress Port or Ingress Port

After configuring a port as a mirroring port, user should use it to monitor the flows of egress or ingress.

Configure Monitor the flows of ingress port:

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input ingress ge-1/1/2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

Configure Monitor the Flows of Egress Port:

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input egress ge-1/1/3
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.  
Save done.
```

Mirroring Configuration Example

As shown in Fig 1, ge-1/1/1 is the ingress port, ge-1/1/2 is the egress port, and ge-1/1/3 is the mirroring port. In the following example, the mirroring port can analyze the flows of egress port and ingress port.

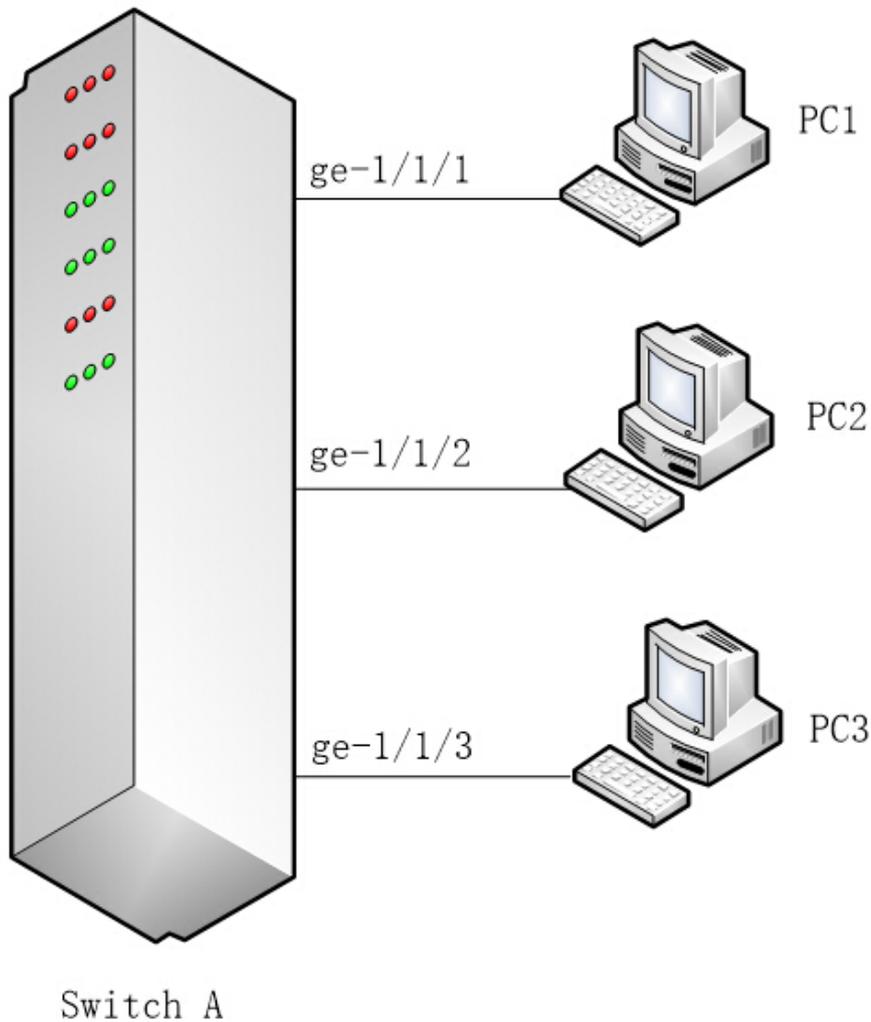


Fig 1. Configure Mirroring

Configure Mirroring Port

Mirroring ports can be used for analyzing the traffic of the egress or ingress port. The following example describes how to configure a mirroring port. In this example, ge-1/1/3 is the mirroring port.

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 output ge-1/1/3  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done
```

Configuring Mirroring on Egress or Ingress Port

The following example shows how to configure mirroring on egress or ingress ports. In this example, the egress port is ge-1/1/2, and the ingress port is ge-1/1/1.

```
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input ingress ge-1/1/1
admin@XorPlus# set interface ethernet-switching-options analyzer 111 input egress ge-1/1/2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

Generate Traffic

PC1 sends traffic to PC2. The expected result is that PC3 can monitor the traffic from the ingress port ge-1/1/1 and the egress port ge-1/1/2.

Note:

1. The mirroring port can belong to any VLAN. This port can be either a trunk port or an access port and will not participate in Layer2 or Layer3 forwarding.
2. The egress port or ingress port can be either an access port or a trunk port.
3. When user sends untagged packets, the priority of mirroring is higher than the priority of adding tag.
4. When user receives tagged packets, the priority of mirroring is higher than the priority of removing tag.
5. The mirroring port can also analyze BPDU/LACP/LLDP packets.
6. When user configures ACL for ingress/egress port, the priority of mirroring is higher than the priority of filter.

BPDU Tunneling Configuration

As a Layer2 tunneling technology, BPDU tunneling enables Layer2 protocol packets from geographically dispersed customer networks to be transparently transmitted over specific tunnels across a service provider network.

Configuring BPDU Tunneling for STP on an Interface

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/37 family ethernet-switching
bpd-tunneling
protocol stp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
```

```
Save done.
admin@XorPlus#
```

Configuring Destination Multicast MAC Address for BPDU Packets

```
admin@XorPlus# set interface bpdu-tunneling destination-mac 01:0E:00:00:00:01
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

BPDU Tunneling Configuration Example

The topology below is an example of BPDU tunneling configuration.

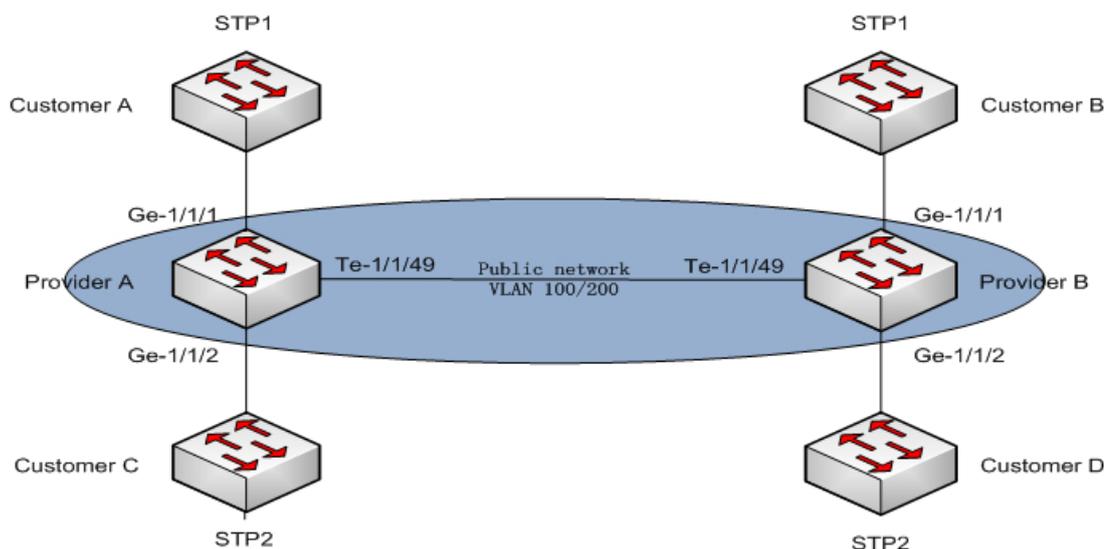


Figure4-12. BPDU Tunneling Configuration.

Configuration on Provider A

Configure VLAN 100 as the default VLAN of Gigabit Ethernet ge-1/1/1, and enable BPDU tunneling on Gigabit Ethernet ge-1/1/1.

```
admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
bpdu-tunneling
protocol stp
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.
Save done.
admin@XorPlus#
```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet ge-1/1/2, and enable BPDU tunneling on Gigabit Ethernet ge-1/1/2.

```
admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
bpdu-tunneling
protocol stp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet te-1/1/49.

```
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

Configure the destination multicast MAC address for BPDUs as 01:0E:00:00:00:1.

```
admin@XorPlus# set interface bpdu-tunneling destination-mac 01:0E:00:00:00:01
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuration on Provider B

Configure VLAN 100 as the default VLAN of Gigabit Ethernet ge-1/1/1, and enable BPDU tunneling on Gigabit Ethernet ge-1/1/1.

```
admin@XorPlus# set vlans vlan-id 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 100
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
```

```

bpd-tunneling
protocol stp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet ge-1/1/2, and enable BPDU tunneling on Gigabit Ethernet ge-1/1/2.

```

admin@XorPlus# set vlans vlan-id 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 200
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
bpd-tunneling
protocol stp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure VLAN 200 as the default VLAN of Gigabit Ethernet te-1/1/49.

```

admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 100
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching vlan
members 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configure the destination multicast MAC address for BPDUs as 01:0E:00:00:00:1.

```

admin@XorPlus# set interface bpd-tunneling destination-mac 01:0E:00:00:00:01
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Flex Links Preemption Delay

User can configure two physical ports or two LAGs as Flex Links or configure one physical port and one LAG as Flex Links.

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 backup-port interface ael
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.

```

```

Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 backup-port delay 10
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface aggregate-ethernet ae2 backup-port interface ae3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configuring the Preemption Mode

By default, the preemption mode is set, and the active interface is preferred. Beyond that, user can configure the "bandwidth" or "off" mode. The "bandwidth" mode calls for a higher bandwidth interface, and the "off" mode turns off preemption.

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 backup-port mode bandwidth
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Showing Flex Links on All Interfaces

User can view the state of user's Flex Links interfaces:

```

admin@XorPlus# run show interface flexlink
Active Interface Backup Interface Mode Delay(seconds)
-----
ge-1/1/1(up) ge-1/1/2(standby) bandwidth 10
admin@XorPlus#

```

Configuring IPv6 RA Guard

When the switch receives an ingress router advertisement (RA) message, it will attempt to match the message via the RA guard. If the ingress port has the RA guard applied but is not a trusted port, the applied VLAN ID will be matched first. If the RA tag is matched with the VLAN ID, the RA guard will continue matching conditions to determine whether to forward or drop the RA message. If the RA tag is *not* matched with the VLAN ID, the applied interface will be matched (followed by the subsequent conditions). RA guard policy can be configured using hop-limit, managed-config-flag, other-config-flag, prefix, source-ipv6-addr, and source-mac-addr options.

```

admin@XorPlus# set protocols neighbour ra-guard 1 hop-limit 1
admin@XorPlus# set protocols neighbour ra-guard 1 managed-config-flag false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols neighbour ra-guard 2 prefix 2001:1:1:1::/64
admin@XorPlus# commit

```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols neighbour ra-guard 3 source-mac-addr 22:22:22:22:22:22
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Trusted-Port

RA guard can be applied to physical interfaces, LAGs, or VLANs. No more than one RA guard can be applied to one interface. The RAs will be forwarded only if all conditions are matched. But if "trusted-port" has been configured for the RA guard, RAs will be forwarded on the trusted port regardless.

```

admin@XorPlus# set protocols neighbour ra-guard term 1 interface ge-1/1/1
admin@XorPlus# set protocols neighbour ra-guard term 1 interface ael
admin@XorPlus# set protocols neighbour ra-guard term 1 vlan-id 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#set protocols neighbour ra-guard trusted-port ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols neighbour ra-guard term 2 vlan-id 3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Displaying RA Guards

```

admin@XorPlus# run show raguard
Raguard: 1
cur hop limit : 1..10
managed configuration : Unset
other configuration : Set
source mac address :
22:22:22:22:22:22
source ipv6 address :
fe80::/64
prefix :
2001:1:1:1::/64
interface : ge-1/1/1, ael
vlan : 2
packet dropped: 0
packet total : 0
Raguard: 2
vlan : 3
packet dropped: 0
packet total : 0
trusted port:
ge-1/1/1
admin@XorPlus#

```

LFS Configuration

Abstract

Link fault signaling (LFS) operates between the remote RS and the local RS. Faults detected between the remote RS and the local RS are received by the local RS as Local Fault. RS is the only layer that can generate Remote Fault signals.

Sublayers within the PHY are capable of detecting faults that render a link unreliable for communication. Upon recognition of a fault condition, a PHY sublayer indicates Local Fault status on the data path. When this Local Fault status reaches an RS, the RS stops sending MAC data, and continuously generates a Remote Fault status on the transmit data path (possibly truncating a MAC frame being transmitted). When Remote Fault status is received by an RS, the RS stops sending MAC data, and continuously generates Idle control characters. When the RS no longer receives fault status messages, it returns to normal operation, sending MAC data.

The RS reports the fault status of the link. Local Fault indicates a fault detected on the receive data path between the remote RS and the local RS. Remote Fault indicates a fault on the transmit path between the local RS and the remote RS.

The fault status is as follows:

a) link_fault = OK

The RS shall send MAC frames as requested through the PLS service interface. In the absence of MAC frames, the RS shall generate Idle control characters.

b) link_fault = Local Fault

The RS shall continuously generate Remote Fault Sequence ordered_sets.

c) link_fault = Remote Fault

The RS shall continuously generate Idle control characters.

Link Fault Signaling

If ignore local fault is set as false: When link local fault is triggered, the RS shall continuously generate Remote Fault Sequence ordered_sets. Otherwise, the RS will not generate Remote Fault Sequence ordered_sets.

If ignore remote fault is set as false: When link remote fault is received, the RS shall continuously generate Idle control characters. Otherwise, the RS shall send MAC frames as requested through the PLS service interface and generate Idle control characters in the absence of MAC frames

LFS Commands

The following is the configuration command as sample:

```
set interface gigabit-ethernet te-1/1/1 link-fault-signaling ignore-local-fault true/false
set interface gigabit-ethernet te-1/1/1 link-fault-signaling ignore-remote-fault true/fals
```

Up Mode

The force up command forcibly brings up a fiber Ethernet port and enables the port to forward packets unidirectionally over a single link. In this way, transmission links are well utilized.

Up Mode Commands

The following is the configuration command as sample:

```
set interface gigabit-ethernet te-1/1/1 up-mode true/false
```



Up-mode true command should be configured with ignore-local-fault true command together. If user only configures up-mode true and does not configure ignore-local-fault command, traffic cannot transmit from TX link.

Symmetric Hash for ECMP and LAG Configuration Example

Symmetric Hash support for ECMP and LAG port. Symmetric hash need the hashing field must be symmetric. For example, packet 1 and packet 2 are symmetric in table 1 and table 2 as below, and then packet 1 and packet2 will go out from the same physical port. Now Our symmetric Hash use IP layer and L4 field to hash when packets are transmitted on ECMP and LAG port. Only matching symmetric condition, two packets can be transmitted on the same member port of LAG interface or ECMP.

IP Packet	Source IP Address	Destination IP Address
Packet1	10.1.1.1	20.1.1.1
Packet2	20.1.1.1	10.1.1.1

Table 1.

Layer 4 Packet	Source IP Address	Destination IP Address	Source Port Number	Destination Port Number
Packet1	10.1.1.1	20.1.1.1	100	200
Packet2	20.1.1.1	10.1.1.1	200	100

Table 2.

Symmetric Hash field as below:

1. ip-source
2. ip-destination
3. port-source
4. port-destination

By default, Enable hash field on LAG interface:

1. ingress-interface
2. ethernet-source-address
3. ethernet-destination-address
4. ethernet-type
5. vlan
6. ip-protocol
7. ip-source
8. ip-destination
9. port-source
10. port-destination

By default, Enable hash field on ECMP:

1. ip-source
2. ip-destination
3. port-source
4. port-destination

 LAG interface enables 10 fields to hash more than symmetric hashing 4 fields in the default case. So if to need symmetric hash work on lag interface please it is best to disable the following field

1. ingress-interface
2. ethernet-source-address
3. ethernet-destination-address
4. ethernet-type
5. vlan
6. ip protocol

 The Symmetric hashing is supported on Trident 2, Trident 2+, Trident 3, Triumph.sdk version need bcm6.3.9 or advanced version.

ECMP Configuration

```
set interface ecmp hash-mapping symmetric true
```

ECMP Examples:

configure three ECMP route:

```
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 199
set interface gigabit-ethernet te-1/1/2 family ethernet-switching native-vlan-id 299
set interface gigabit-ethernet te-1/1/3 family ethernet-switching native-vlan-id 300
set interface gigabit-ethernet te-1/1/4 family ethernet-switching native-vlan-id 301
set protocols static route 100.100.100.0/24 next-hop 182.168.1.100
set protocols static route 100.100.100.0/24 qualified-next-hop 183.168.1.100 metric 1
set protocols static route 100.100.100.0/24 qualified-next-hop 184.168.1.100 metric 1
set protocols static route 172.168.1.0/24 next-hop 182.168.1.100
```

```

set protocols static route 172.168.1.0/24 qualified-next-hop 183.168.1.100 metric 1
set protocols static route 172.168.1.0/24 qualified-next-hop 184.168.1.100 metric 1
set vlan-interface interface vlan199 vif vlan199 address 172.168.1.1 prefix-length 24
set vlan-interface interface vlan299 vif vlan299 address 182.168.1.1 prefix-length 24
set vlan-interface interface vlan300 vif vlan300 address 183.168.1.1 prefix-length 24
set vlan-interface interface vlan301 vif vlan301 address 184.168.1.1 prefix-length 24
set vlan-interface interface vlan399 vif vlan399 address 100.100.100.1 prefix-length 24
set vlans vlan-id 199 l3-interface "vlan199"
set vlans vlan-id 299 l3-interface "vlan299"
set vlans vlan-id 300 l3-interface "vlan300"
set vlans vlan-id 301 l3-interface "vlan301"
set vlans vlan-id 399 l3-interface "vlan399"

```

Configure the symmetric hash true

```

set interface ecmp hash-mapping symmetric true

```

LAG Hashing Configuration

```

set interface aggregate-ethernet ael hash-mapping mode advanced
set interface aggregate-balancing hash-mapping symmetric true

```

LAG Hashing Examples:

configure one lag with three ports:

```

set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 199
set interface aggregate-ethernet ael family ethernet-switching native-vlan-id 299
set interface gigabit-ethernet te-1/1/2 ether-options 802.3ad ael
set interface gigabit-ethernet te-1/1/3 ether-options 802.3ad ael
set interface gigabit-ethernet te-1/1/4 ether-options 802.3ad ael
set protocols static route 100.100.100.0/24 next-hop 182.168.1.100
set protocols static route 172.168.1.0/24 next-hop 182.168.1.100
set vlan-interface interface vlan199 vif vlan199 address 172.168.1.1 prefix-length 24
set vlan-interface interface vlan299 vif vlan299 address 182.168.1.1 prefix-length 24
set vlan-interface interface vlan399 vif vlan399 address 100.100.100.1 prefix-length 24
set vlans vlan-id 199 l3-interface vlan199
set vlans vlan-id 299 l3-interface vlan299
set vlans vlan-id 399 l3-interface vlan399

```

Configure the symmetric hash true

```

set interface aggregate-ethernet ael hash-mapping mode advanced
set interface aggregate-balancing hash-mapping symmetric true
set interface aggregate-balancing hash-mapping field ingress-interface disable true
set interface aggregate-balancing hash-mapping field ethernet-source-address disable true
set interface aggregate-balancing hash-mapping field ethernet-destination-address disable true
set interface aggregate-balancing hash-mapping field ethernet-type disable true
set interface aggregate-balancing hash-mapping field vlan disable true
set interface aggregate-balancing hash-mapping field ip-protocol disable true

```

Layer 3 Unicast Routing Configuration

This chapter describes the configuration steps of Layer 3 routing, including static routing, RIPv2, OSPFv2, VRRP, and ECMP.

- Layer 3 VLAN Interface Configuration
- ARP Configuration
- Dynamic ARP Inspection
- Default Administrative Distance Values
- Static Routing Configuration
- Static Routing Configuration Example
- RIPv2 Routing Protocol Configuration
- RIPv2 Routing Configuration Example
- OSPF (Open Shortest Path First)
 - OSPF Overview
 - Basic OSPF Configuration Tasks
 - Basic OSPF Configuration Example
- OSPF Configuration Example: NSSA Stub Normal
- OSPF Stub Area NSSA Summary
- OSPF Virtual Link Configuration Guide
- OSPF Area Range Configuration Guide
- Importing an External Route into an OSPF Area
- BFD Protocol Configuration
- BFD Basic Configuration Example
- Configuring ECMP (Equal-Cost Multipath Routing)
- Configuring VRRP (Virtual Router Redundancy Protocol)
- Configuring Active-Active-VRRP(Active Active Virtual Router Redundancy Protocol)
- IPv6 Neighbor Configuration
- IPv6 Static Routing Configuration
- OSPFv3 Routing Protocol Configuration
- ACL and Filter Configuration
 - Configuring Control Plane Security Guide
 - Configuration Control Plane Security Policer
 - Control Plane Security Examples

- IPv4/IPv6 BGP Configuration
 - IPv4 BGP configuration
 - BGP Configuration Guide
 - BGP Basic Configuration Example
 - BGP Route Reflector Configuration Example
 - BGP Confederation Configuration Example
 - BGP Load Balancing Configuration Example
 - IPv6 BGP Configuration
 - IPv6 BGP Introduction
 - Building Peering Sessions
 - EBGP Peering
 - IBGP Peering
 - Establish BGP Peer Use 4-byte-AS-Number
 - Sources of Routing Updates
 - Injecting Information Dynamically into BGP
 - Injecting Information Statically into BGP
 - BGP Attributes
 - The NEXT_HOP Attribute
 - The AS_PATH Attribute
 - The LOCAL_PREF Attribute
 - The MULTI_EXIT_DISC Attribute
 - The COMMUNITY Attribute
 - BGP-4 Aggregation
 - Synchronization
 - Controlling Large-Scale Autonomous System
 - Confederations
 - Route Reflectors
 - Redundancy and Load Balancing
 - Designing Stable Internets
 - Label BGP
 - Labeled BGP Support
 - Configuration Example for Labeled Support
 - IPV4 Labeled BGP Configuration
 - IPV6 Labeled-BGP Configuration

- Debugging CLI for Labeled-BGP

Layer 3 VLAN Interface Configuration

The Layer 3 interface *is* a VLAN interface. Use the following configuration sequence:

1. Create a VLAN using the `set vlans vlan-id` command.
2. Associate an L3 interface with the VLAN using the `vlans vlan-id l3-interface` command.
3. Configure the IP address and prefix length for the virtual interface using the `vlan-interface interface vif address prefix-length` command.



When all the interfaces in a VLAN are link-down, the VLAN interface will be link-down. The VLAN interface will be link-up only when at least one of the member interfaces is link-up.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show vlan-interface
vlan-2 Hwaddr C8:0A:A9:9E:14:9F, Vlan:2, State:DOWN
Inet addr: 192.168.1.1/24
fe80::ca0a:a9ff:fe9e:149f/64
Traffic statistics:
IPv4 Input Packets.....0
IPv4 Forwarding Packets.....0
IPv6 Input Packets.....0
IPv6 Forwarding Packets.....0
vlan-3 Hwaddr C8:0A:A9:9E:14:9F, Vlan:3, State:UP
Inet addr: 192.168.2.1/24
fe80::ca0a:a9ff:fe9e:149f/64
Traffic statistics:
IPv4 Input Packets.....0
IPv4 Forwarding Packets.....0
IPv6 Input Packets.....0
IPv6 Forwarding Packets.....0
admin@XorPlus#
```

ARP Configuration

Configuring ARP Aging Time

In the default setting, the ARP aging time is 1200 seconds.

```
admin@XorPlus# set protocols arp aging-time 600
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring a Static ARP Entry

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlan-interface interface vlan-2 address 192.168.1.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#set protocols arp interface vlan-2 address 192.168.1.1 mac-address
22:22:22:22:22:22
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Dynamic ARP Inspection

Dynamic ARP Inspection (DAI), is a security feature that validates ARP packets in a network. DAI intercepts and discards ARP packets with invalid IP-to-MAC address bindings. This capability protects the network from some man-in-the-middle attacks. DAI ensures that only valid ARP requests and responses are relayed. The switch performs these activities:

- Intercepts all ARP requests and responses on untrusted ports.
- Verifies that each of these intercepted packets has a valid IP-to-MAC address binding before updating the local ARP cache, or before forwarding the packet to the appropriate destination.
- Drops invalid ARP packets.

DAI determines the validity of an ARP packet based on valid IP-to-MAC address bindings stored in a trusted database, the DHCP snooping binding database. This database is built by DHCP snooping, if DHCP snooping is enabled on the VLANs and on the switch. If the ARP packet is received on a trusted interface, the switch forwards the packet without any checks. On untrusted interfaces, the switch forwards the packet only if it is valid.

DAI associates a trust state with each interface on the switch. Packets arriving on trusted interfaces bypass all DAI validation checks, and those arriving on untrusted interfaces undergo the DAI validation process.

In a typical network configuration, you configure all switch ports connected to host ports as untrusted, and configure all switch ports connected to switches as trusted. With this configuration, all ARP packets entering the network from a given switch bypass the security check. No other validation is needed at any other place in the VLAN or in the network.

When configuring DAI, follow these guidelines and restrictions:

- DAI is an ingress security feature; it does not perform any egress checking.
- DAI is not effective for hosts connected to switches that do not support DAI, or that do not have this feature enabled. Because man-in-the-middle attacks are limited to a single Layer 2 broadcast domain, separate the domain with DAI checks from the one with no checking. This action secures the ARP caches of hosts in the domain enabled for DAI.
- DAI depends on the entries in the DHCP snooping binding database to verify IP-to-MAC address bindings in incoming ARP requests and ARP responses. Make sure to enable DHCP snooping to permit ARP packets that have dynamically assigned IP addresses.
- DAI is supported on access ports, trunk ports, and lag ports.

DAI Configuration example

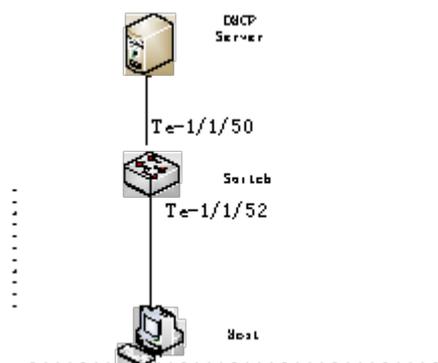


Figure-DAI

Step 1: Enable DHCP snooping on Switch.

You can enable DHCP snooping on the egress port, the port connected to DHCP server.

Enable DHCP snooping:

```
admin@XorPlus# set protocols dhcp snooping disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
```

Set the interface to trust mode

```
admin@XorPlus# set protocols dhcp snooping port te-1/1/50 trust true
admin@XorPlus# commit
Commit OK.
```

```
Save done.
admin@XorPlus#
```

Step 2: Enable DAI

You can enable DAI on the port connect to the host.

```
admin@XorPlus# set protocols arp interface vlan-900 inspection disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
admin@XorPlus#
```

Step 3: Check ARP inspection table.

When the host gets an IP address from the DHCP server, and the switch has enabled DHCP snooping, it will create a table, the IP-MAC-port binded table. The entry in this table was trusted, all other ARP packets not in this table, will be discarded. (The ARP packet must be in accordance with the ARP inspection table, interface, IP address, and MAC address must be identified).

```
admin@XorPlus# run show arp inspection
Total count : 1
Interface DAI Address HW Address
-----
vlan-900 Enabled 192.168.9.5 0:1e:c9:bb:d3:35
```

Default Administrative Distance Values

Select the Best Path

A PicOS system can run multiple routing protocols simultaneously. For example, RIP may be used to distribute routes within our network and BGP to learn external routes. In some situations this can lead to a router learning the same route from more than one routing protocol.

For example, these two routes might be learned:

- Subnet: 128.16.64.0/24, nexthop: 192.150.187.1, learned from BGP via an external peering. AS Path: 123 567 987.
- Subnet: 128.16.64.0/24, nexthop: 10.0.0.2, learned from RIP with metric 13

The longest prefix match rule doesn't help us because the prefix lengths are the same, and the metric used for RIP is not directly comparable against the AS path length or any other attribute attached to a BGP route. So, how do we decide which route to take? A PicOS system uses the concept of administrative distance to determine which route wins. Basically, each routing protocol has a configured "distance," and if a route is heard from two protocols, then the version with the smallest distance wins.

Default Distance Value Table

This table lists the administrative distance default values of the protocols that PicOS supports:

Route Protocol	Default Distance Values
Directly connected	0
Static route	1
External BGP	20
OSPF	110
RIP	120
Internal BGP	200

Hence, in the example above, the route learned from BGP will be preferred.



The administrative distance is fixed in the PicOS system, and cannot be modified.

Static Routing Configuration

In L2/L3, all routing entries will be configured to the ASIC switching chip if the outgoing VLAN-interface is link-up, and the outgoing physical port is learning.

Traffic that can be routed will have a route entry in the RIB and the ARP of the next hop; the outgoing interface should be link-up. The traffic will then be soft-routed (i.e., routed by the switch's CPU).

When the switch learns the MAC address of the next-hop, the switch will forward the traffic with the ASIC chip.

Configure static route:

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-2 address 192.168.2.1
prefix-length 24
admin@XorPlus# set protocols static route 10.10.1.0/24 next-hop 192.168.2.5
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show route table ipv4 unicast final
10.10.1.0/24 [static(1)/1]
> to 192.168.2.5 viavlan-3/vlan-3
```

```

192.168.1.0/24 [connected(0)/0]
> via vlan-2/vlan-2
192.168.2.0/24 [connected(0)/0]
> via vlan-3/vlan-3
admin@XorPlus#
admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
10.10.1.0 255.255.255.0 00:1E:68:37:EF:7D ge-1/1/2
192.168.1.0 255.255.255.0 C8:0A:A9:04:49:28 connected
192.168.2.0 255.255.255.0 C8:0A:A9:04:49:28 connected

```

With the **show route forward-routeipv4 all** command, all the route entries in the ASIC chip will be displayed. Following the **show route table ipv4 unicastfinal** command, all routes in the RIB of the kernel will be displayed.

Configure max-route-limit.

Before configuring max-route-limit, check the forward-route table:

```

admin@XorPlus# run show route forward-route ipv4 all
Destination      NetMask          NextHopMac       Port
-----
192.168.1.0      255.255.255.0   04:7D:7B:62:93:FF connected
192.168.2.0      255.255.255.0   04:7D:7B:62:93:FF connected
192.168.3.0      255.255.255.0   04:7D:7B:62:93:FF connected
192.168.20.0     255.255.255.0   04:7D:7B:62:93:FF connected
192.168.100.0    255.255.255.0   04:7D:7B:62:93:FF connected
Total route count:5
admin@XorPlus#
admin@XorPlus# run show route forward-route ipv6 all
Destination      NetMask
NextHopMac       Port
-----
5001::          ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
2001::          ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
7001::          ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 te-1/1/3
6001::          ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 te-1/1/3
Total route count:4
admin@XorPlus#

```

After configuring max-route-limit:

```

admin@XorPlus# set interface max-route-limit 1
admin@XorPlus# comm
Merging the configuration.
Commit OK.
Save done.
Maximum of route limit changes, please reboot system to make it effect!
admin@XorPlus#

```

Check the forward-route table:

```

admin@XorPlus# run show route forward-route ipv4 all
Destination      NetMask          NextHopMac       Port
-----
192.168.1.0      255.255.255.0   04:7D:7B:62:93:FF connected

```

```

192.168.2.0      255.255.255.0    04:7D:7B:62:93:FF  connected
192.168.3.0      255.255.255.0    04:7D:7B:62:93:FF  connected
192.168.20.0     255.255.255.0   04:7D:7B:62:93:FF  connected
192.168.100.0    255.255.255.0   04:7D:7B:62:93:FF  connected
Total route count:5
admin@XorPlus#
admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                                Port
-----
5001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF  connected
2001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF  connected
Total route count:2
admin@XorPlus#

```

Note: The switch supports 12k ipv4 routes at most and 6k ipv6 routes at most. One ipv6 route is equal to two ipv4 routes. The directly-connected routes are not excepted to route limit.

Configuration Example

An example of configuration with static routing is shown in Fig. 5-1. Host A and Host B should be able to communicate with each other. Host A and Host B should be able to communicate with the gateway (e.g., access Internet).

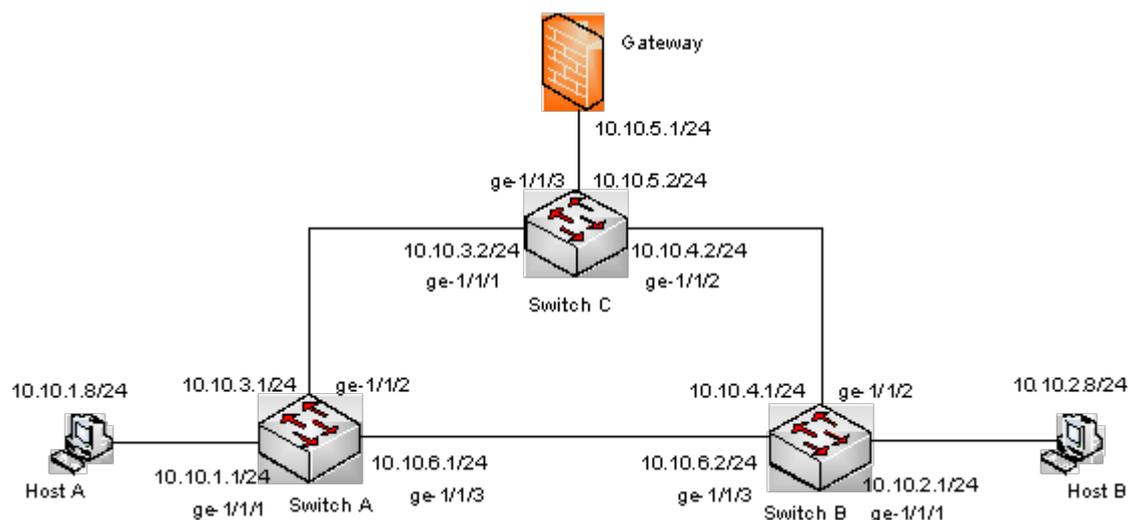


Figure 5-1. Static routing configuration.

Configuring Switch A

For Switch A, configure 3 VLAN interfaces for networks 10.10.1.1/24, 10.10.3.1/24, and 10.10.6.1/24. Also configure a static route to 10.10.2.0/24 and a default route.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2

```

```

admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 4
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlans vlan-id 4 l3-interface vlan-4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-2 address 10.10.3.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-4vif vlan-2 address 10.10.6.1
prefix-length 24
admin@XorPlus# set protocols static route 10.10.2.0/24 next-hop 10.10.6.2
admin@XorPlus# set protocols static route 0.0.0.0/0 next-hop 10.10.3.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verify the route entry in the RIB as follows:

```

admin@XorPlus# run show route table ipv4 unicast final
0.0.0.0/0 [static(1)/1]
> to 10.10.3.2 via vlan-3/vlan-3
10.10.2.0/24[static(1)/1]
> to 10.10.6.2 via vlan-4/vlan-4
10.10.1.0/24 [connected(0)/0]
> via vlan-2/vlan-2
10.10.3.0/24 [connected(0)/0]
> via vlan-3/vlan-3
10.10.6.0/24 [connected(0)/0]
> via vlan-4/vlan-4
admin@XorPlus#

```

Configuring Switch B

Configure 3 VLAN interfaces for networks 10.10.2.1/24, 10.10.4.1/24, and 10.10.6.2/24. Then, configure a static route to 10.10.1.0/24 and a default route.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 4
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlans vlan-id 4 l3-interface vlan-4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

```

admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-2 address 10.10.4.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-4vif vlan-2 address 10.10.6.1
prefix-length 24
admin@XorPlus# set protocols static route 10.10.1.0/24 next-hop 10.10.6.1
admin@XorPlus# set protocols static route 0.0.0.0/0 next-hop 10.10.4.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verify the route entry in the RIB:

```

admin@XorPlus# run show route table ipv4 unicast final
0.0.0.0/0 [static(1)/1]
> to 10.10.4.2 via vlan-3/vlan-3
10.10.1.0/24[static(1)/1]
> to 10.10.6.1 via vlan-4/vlan-4
10.10.2.0/24 [connected(0)/0]
> via vlan-2/vlan-2
10.10.4.0/24 [connected(0)/0]
> via vlan-3/vlan-3
10.10.6.0/24 [connected(0)/0]
> via vlan-4/vlan-4
admin@XorPlus#

```

Configuring Switch C

Configure 3 VLAN interfaces for networks 10.10.3.2/24, 10.10.4.2/24, and 10.10.5.2/24. Then, configure a static route to 10.10.1.0/24 and a default route.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 4
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlans vlan-id 4 l3-interface vlan-4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address 10.10.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address 10.10.4.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-4vif vlan-4 address 10.10.5.2
prefix-length 24
admin@XorPlus# set protocols static route 10.10.1.0/24 next-hop 10.10.3.1
admin@XorPlus# set protocols static route 10.10.2.0/24 next-hop 10.10.4.1
admin@XorPlus# set protocols static route 10.10.6.0/24 next-hop 10.10.3.1
admin@XorPlus# set protocols static route 0.0.0.0/0 next-hop 10.10.5.1
admin@XorPlus# commit

```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verify the route entry in the RIB:

```

admin@XorPlus# run show route table ipv4 unicast final
0.0.0.0/0 [static(1)/1]
> to 10.10.5.1 via vlan-4/vlan-4
10.10.1.0/24[static(1)/1]
> to 10.10.3.1 via vlan-2/vlan-2
10.10.2.0/24[static(1)/1]
> to 10.10.4.1 via vlan-3/vlan-3
10.10.6.0/24[static(1)/1]
> to 10.10.3.1 via vlan-2/vlan-2
10.10.3.0/24 [connected(0)/0]
> via vlan-2/vlan-2
10.10.4.0/24 [connected(0)/0]
> via vlan-3/vlan-3
10.10.5.0/24 [connected(0)/0]
> via vlan-4/vlan-4
admin@XorPlus#

```

RIPv2 Routing Protocol Configuration

In L2/L3, RIPv2 is supported. A policy statement is used to specify which route entry will be distributed. For example, static route or the connected route can be distributed to a neighbor. Also, the distributed route metric can be specified. The RIPv2 interface parameters (accept-default-route, advertise-default-route, deletion-delay, request-interval, update-interval) can be configured.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1
prefix-length 24
admin@XorPlus# set protocols static route 9.9.9.0/24 next-hop 192.168.2.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set policy policy-statement connected-to-rip term export from protocol
connected
admin@XorPlus# set policy policy-statement connected-to-rip term export then metric 0
admin@XorPlus# set policy policy-statement static-to-rip term export from protocol static
admin@XorPlus# set policy policy-statement static-to-rip term export then metric 1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols rip interface vlan-2 vif vlan-2 address 192.168.1.1

```

```

admin@XorPlus# set protocols rip export "connected-to-rip,static-to-rip"
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

You can verify the RIP configuration:

```

admin@XorPlus# run show rip status all
* RIP on vlan-2vlan-2 192.168.1.1
  Status: enabled
admin@XorPlus#
admin@XorPlus# run show rip statistics all
* RIP on vlan-2vlan-2 192.168.1.1
  Status: enabled
  Counter                                     Value
-----
Requests Sent                               7
Updates Sent                                6
Triggered Updates Sent                      1
Non-RIP Updates Sent                        0
Total Packets Received                      0
Request Packets Received                    0
Update Packets Received                     0
Bad Packets Received                        0
Authentication Failures                     0
Bad Routes Received                         0
Non-RIP Requests Received                   0

```

RIPv2 Routing Configuration Example

Configuration Example

An example of configuring RIPv2 is shown in Fig. 5-2. Host A and Host B should be able to communicate with each other with an RIP route. Host A and Host B should be able to communicate with the gateway (e.g., access Internet) with RIP.

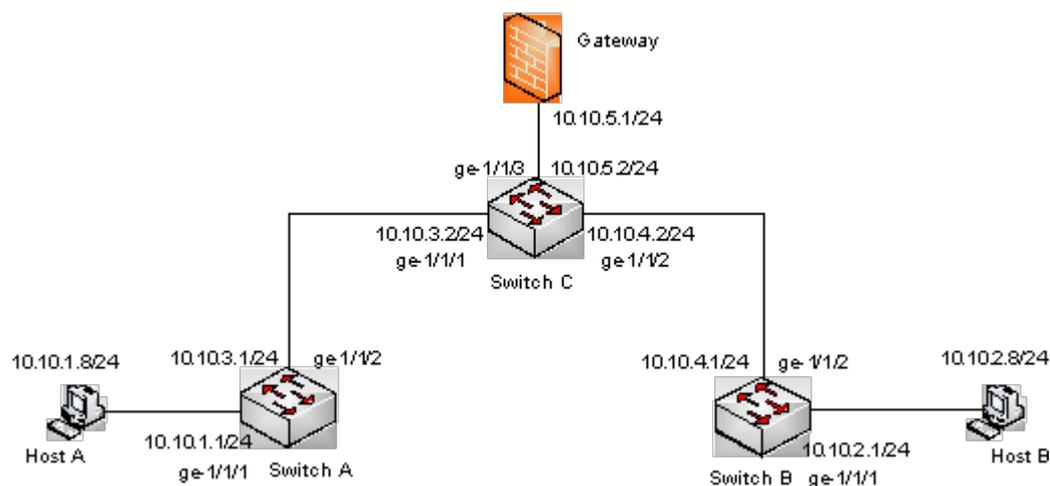


Figure 5-2. RIPv2 routing configuration.

Configuring Switch A

For Switch A, configure 2 VLAN interfaces for networks 10.10.1.1/24 and 10.10.3.1/24. Then, configure an RIP interface in network 10.10.3.1/24. Switch A should accept the default route, which is advertised by Switch C.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address 10.10.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address 10.10.3.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set policy policy-statement connected-to-rip term export from protocol
connected
admin@XorPlus# set policy policy-statement connected-to-rip term export then metric 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols rip interface vlan-3 address 10.10.3.1
admin@XorPlus# set protocols rip export "connected-to-rip"
admin@XorPlus# set protocols rip interface vlan-3vif vlan-3 address 10.10.3.1
accept-default-route true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B

Configure 2 VLAN interfaces for networks 10.10.1.1/24 and 10.10.3.1/24. Then, configure an RIP interface in network 10.10.3.1/24. Switch B should accept the default route, which is advertised by Switch C.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address 10.10.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-2 address 10.10.4.1
prefix-length 24
admin@XorPlus# commit
```

```

Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set policy policy-statement connected-to-rip term export from protocol
connected
admin@XorPlus# set policy policy-statement connected-to-rip term export then metric 0
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols rip interface vlan-3vif vlan-3 address 10.10.4.1
admin@XorPlus# set protocols rip export "connected-to-rip"
admin@XorPlus# set protocols rip interface vlan-3vif vlan-3 address 10.10.4.1
accept-default-route true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure 3 VLAN interfaces for networks 10.10.3.2/24, 10.10.4.2/24, and 10.10.5.2/24. Then, configure a default route and 2 RIP interfaces.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 4
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlans vlan-id 4 l3-interface vlan-4
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address 10.10.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address 10.10.4.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-4vif vlan-4 address 10.10.5.2
prefix-length 24
admin@XorPlus# set protocols static route 0.0.0.0/0 next-hop 10.10.5.1
admin@XorPlus# set protocols rip interface vlan-2vif vlan-2 address 10.10.3.2
admin@XorPlus# set protocols rip interface vlan-2vif vlan-2 address 10.10.3.2
advertise-default-route true
admin@XorPlus# set protocols rip interface vlan-3vif vlan-3 address 10.10.4.2
admin@XorPlus# set protocols rip interface vlan-3vif vlan-3 address 10.10.4.2
advertise-default-route true
admin@XorPlus# set protocols rip export "connected-to-rip"
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verifying the RIP Configuration

Verify the RIP configuration of the switches as shown below. (In the following example, the RIP peer and the RIP route table in Switch A is verified.)

```
admin@XorPlus# run show rip peer
Address Interface State Hello Rx Hello Tx Last Hello
-----
10.10.3.2vlan-3/vlan-3 Up 0 0 00:41:44
admin@XorPlus#
admin@XorPlus# run show route table ipv4 unicast rip
0.0.0.0/0[rip(120)/1]
> to 10.10.3.2 via vlan-3/vlan-3
10.10.2.0/24 [rip(120)/1]
> to 10.10.3.2 via vlan-3/vlan-3
10.10.4.0/24 [rip(120)/1]
> to 10.10.3.2 via vlan-3/vlan-3
```

OSPF (Open Shortest Path First)

OSPF (Open Shortest Path First) is a link-state routing protocol that is not proprietary to any vendor or organization. OSPF was developed by IETF (Internet Engineering Task Force), and it is the IGP (Interior Gateway Protocol) recommended by Pica8.

OSPF has evolved through several RFCs. OSPF version 2, which is the current version for IPv4, is defined in RFC 2328. PicOS supports version 2 of the OSPF protocol in L2/L3 mode.

Pica8 PicOS Version 2.6 supports the following OSPF area types:

1. Normal Area
2. Stub Area
3. NSSA (Not-So-Stubby Area)
 - OSPF Overview
 - Basic OSPF Configuration Tasks
 - Basic OSPF Configuration Example

OSPF Overview

The high-level operation of OSPF is explained below:

1. OSPF routers send Hello packets out of all OSPF-enabled interfaces. Two routers sharing a common data link become *neighbors*. If they agree on certain parameters in Hello packets.
2. Some neighbors move on and form *adjacencies*, which can be thought of as virtual point-to-point links over which routing information is exchanged.
3. Each OSPF router sends LSAs (link-state advertisements) over all its adjacencies. The LSAs describe the router's neighbors, links, and the state of the links. OSPF defines multiple LSA types to communicate different types of link-state information.

4. When an OSPF router receives an LSA from a neighbor, it adds the LSA to its link-state database. The router also sends a copy of the LSA over all of its adjacencies. The flooding of LSAs throughout an OSPF area enables all routers to have identical link-state databases.
5. When the link-state databases are built, every router runs the Dijkstra's SPF (Shortest Path First) algorithm to calculate the shortest loop-free path to every known subnet. The collection of all paths calculated by the router, with itself as the root, is known as the SPF tree.
6. Each router populates its routing table from its SPF tree.

Basic OSPF Configuration Tasks

To configure OSPF on a PicOS device, complete the tasks described in the following sections .

Configuring OSPF Router ID

The OSPF router ID is a 32-bit value similar to an IP address by which the OSPF router or switch is uniquely identified within the OSPF domain.

Use the following command to configure OSPF router ID:

```
set protocols ospf4 router-id <router-id>
```

The following example configures 1.1.1.1 as the OSPF router ID:

```
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# commit
Commit OK.
Save done.
```

Do not change the router ID after completing the configuration.

Configuring OSPF Areas

Divide an OSPF domain into areas, which contains the flow of most routing protocol traffic within a single area and reduces the impact of protocol on CPU and memory.

Use the following command to configure OSPF areas:

```
set protocols ospf4 area <area-id> area-type <type>
```

Area ID 0.0.0.0 is reserved for the backbone and each OSPF domain should have the backbone area. All traffic between two non-backbone areas must pass through the backbone. OSPF area types supported by PicOS are: normal, stub, and NSSA.

The example that follows demonstrates configuration of three different areas:

Area ID	Area Type
0.0.0.0	Normal
1.1.1.1	Stub
2.2.2.2	NSSA

```

admin@XorPlus# set protocols ospf4 area 0.0.0.0 area-type normal
admin@XorPlus# set protocols ospf4 area 1.1.1.1 area-type stub
admin@XorPlus# set protocols ospf4 area 2.2.2.2 area-type nssa
admin@XorPlus# commit
Commit OK.
Save done.

```

Configuring OSPF Interfaces

After configuring an OSPF area, assign a switch interface to the area. Once an interface is assigned to an OSPF area, it will start sending OSPF hello messages.

Use the following command to assign a Layer 3 interface to an OSPF area:

```
set protocols ospf4 area <area-id> interface <interface-name> vif <interface-name> address <ip-address>
```

The following commands create Layer 3 VLAN interfaces vlan-2 and vlan-3. The interfaces are given IP addresses and switch ports are assigned to them. Finally, both interfaces are configured to be in the OSPF backbone area according to the table shown below.

VLAN Interface	IP Address / Mask Length	Assigned Switch Interface(s)	OSPF Area
vlan-2	10.10.60.10 / 24	ge-1/1/1	0.0.0.0
vlan-3	10.10.61.10 / 24	ge-1/1/2	0.0.0.0

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.60.10
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.61.10
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.60.10
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-3 vif vlan-3 address
10.10.61.10
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Additional OSPF Parameters

Fine-tune OSPF operation by configuring additional OSPF interface parameters including hello interval, interface cost, passive interface, priority, retransmit interval, router dead interval, and transmit delay.

The following example demonstrates how to configure some of the OSPF interface parameters:

```

admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.60.10 hello-interval 5
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.60.10 interface-cost 8
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.60.10 transmit-delay 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Basic OSPF Configuration Example

Fig.5-3 presents an example of configuring OSPF routing. Switch A and Switch B are located in the backbone area, 0.0.0.0. There are two non-backbone areas, 1.1.1.1 and 2.2.2.2.

Switch D will obtain the routes of networks 10.10.1.0/24, 10.10.3.0/24, and 10.10.9.0/24 through the LSAs sent from its neighbors. Switch C will obtain the routes of networks 10.10.1.0/24, 10.10.2.0/24, and 10.10.8.0/24 according to LSAs sent from its neighbors.

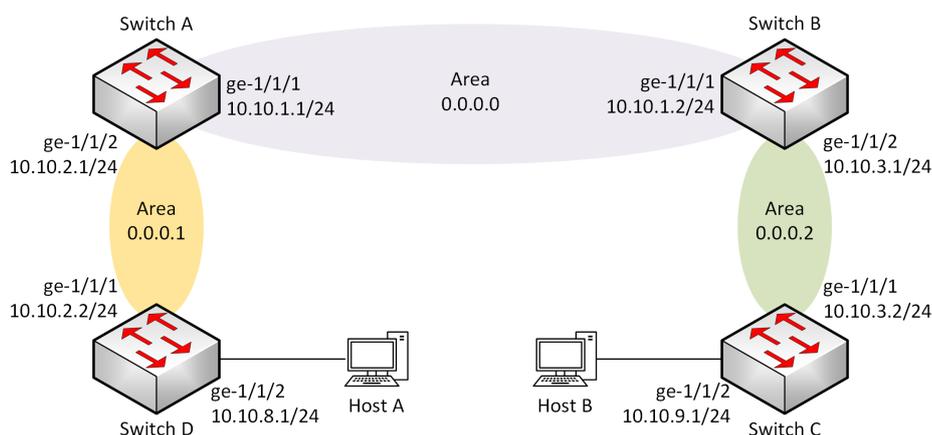


Figure 5-3 Basic OSPF Configuration

Configuring Switch A

For switch A, configure 2 VLAN interfaces for networks 10.10.1.1/24 and 10.10.2.1/24. Also configure area 0.0.0.0, which includes network 10.10.1.1/24, and area 0.0.0.1, which includes network 10.10.2.1/24.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.1

```

```

prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.1 interface vlan-3 vif vlan-3 address
10.10.2.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B

Configure 2 VLAN interfaces for networks 10.10.1.2/24 and 10.10.3.1/24. Then configure area 0.0.0.0, which includes network 10.10.1.2/24, and area 0.0.0.3, which includes network 10.10.3.1/24.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.3.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 2.2.2.2
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.1.2
admin@XorPlus# set protocols ospf4 area 0.0.0.2 interface vlan-3 vif vlan-3 address
10.10.3.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure just one OSPF interface in area 0.0.0.2.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3

```

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.9.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 0.0.0.2 interface vlan-2 vif vlan-2 address
10.10.3.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch D

Configure just one OSPF interface in area 0.0.0.1.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.2.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.8.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4
admin@XorPlus# set protocols ospf4 area 0.0.0.1 interface vlan-2 vif vlan-2 address
10.10.2.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verifying the OSPF configuration

You can verify the OSPF configuration of a switch by checking its OSPF neighbor.

Below, switch A has two OSPF neighbor interfaces, 10.10.1.2 and 10.10.2.2.

```

admin@XorPlus# run show ospf4 neighbor
Address Interface State Router ID Pri Dead
-----
10.10.1.2vlan-2/vlan-2 Full 2.2.2.2 1 32
10.10.2.2vlan-3/vlan-3 Full 4.4.4.4 1 32

```

Then check the OSPF database as shown below:

```

admin@XorPlus# run show ospf4 database
OSPF link state database, Area 0.0.0.0
Type ID Adv Rtr Seq Age Opt Cksum Len
-----
Router *1.1.1.1 1.1.1.1 0x8000025a 394 0x2 0xf2bb 48
Network *10.10.1.1 1.1.1.1 0x80000180 394 0x2 0xc0b9 32
Network *10.10.2.1 1.1.1.1 0x80000180 394 0x2 0xc0b9 32
Router 2.2.2.2 2.2.2.2 0x8000023e 339 0x2 0x3024 36
Network 10.10.3.1 2.2.2.2 0x80000180 394 0x2 0xc0b9 32
Router 3.3.3.3 3.3.3.3 0x8000023e 339 0x2 0x3024 36
Network 10.10.9.1 3.3.3.3 0x80000180 394 0x2 0xc0b9 32
Router 4.4.4.4 4.4.4.4 0x8000023e 339 0x2 0x3024 36
Network 10.10.8.1 4.4.4.4 0x80000180 394 0x2 0xc0b9 32
OSPF link state database, Area 0.0.0.2
Type ID Adv Rtr Seq Age Opt Cksum Len
-----
Router *1.1.1.1 1.1.1.1 0x8000025a 394 0x2 0xf2bb 48
Network *10.10.1.1 1.1.1.1 0x80000180 394 0x2 0xc0b9 32
Network *10.10.2.1 1.1.1.1 0x80000180 394 0x2 0xc0b9 32
Router 2.2.2.2 2.2.2.2 0x8000023e 339 0x2 0x3024 36
Network 10.10.3.1 2.2.2.2 0x80000180 394 0x2 0xc0b9 32
Router 3.3.3.3 3.3.3.3 0x8000023e 339 0x2 0x3024 36
Network 10.10.9.1 3.3.3.3 0x80000180 394 0x2 0xc0b9 32
Router 4.4.4.4 4.4.4.4 0x8000023e 339 0x2 0x3024 36
Network 10.10.8.1 4.4.4.4 0x80000180 394 0x2 0xc0b9 32
Finally, you can check the OSPF route in the RIB of switch A.
admin@XorPlus#
admin@XorPlus# run show route table ipv4 unicast ospf
10.10.3.0/24 [ospf(110)/2]
> to 10.10.1.2 via vlan-2/vlan-2

```

Configuration Example

The configurations of an OSPF NSSA and a stub area are shown in Fig. 5-4. Switch D will obtain the routes of networks 10.10.1.0/24, 10.10.3.0/24, and 10.10.9.0/24 according to the LSAs received from its neighbors. Switch C will obtain the routes of networks 10.10.1.0/24, 10.10.2.0/24, and 10.10.8.0/24 according to the LSAs received from its neighbors. The figure below does not include RIP or BGP configurations.

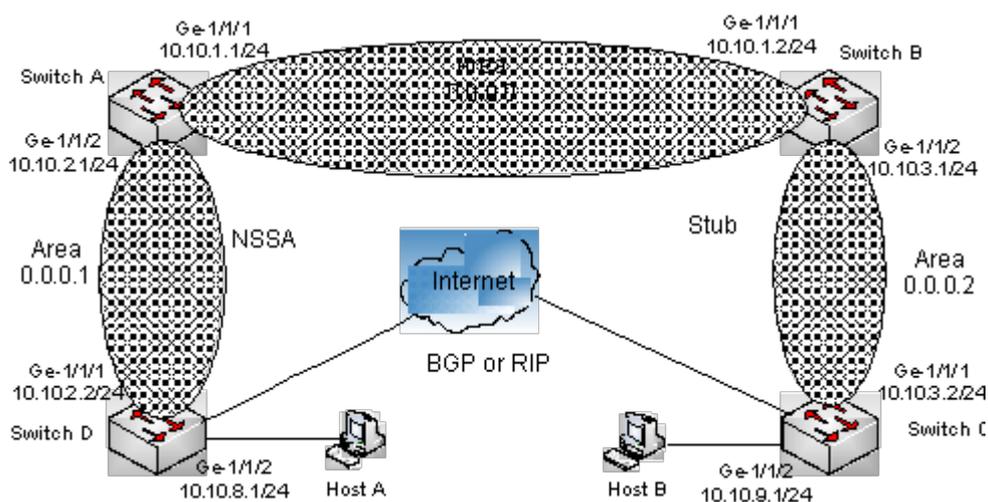


Figure 5-4. OSPF NSSA, stub area configurations.

Configuring Switch A

For switch A, configure 2 VLAN interfaces for networks 10.10.1.1/24 and 10.10.2.1/24. Also configure area 0.0.0.0, which includes network 10.10.1.1/24 and area 0.0.0.1, which includes network 10.10.2.1/24.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.1 interface vlan-3 vif vlan-3 address
10.10.2.1
admin@XorPlus# set protocols ospf4 area 0.0.0.1 area-type nssa
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B

Configure 2 VLAN interfaces for networks 10.10.1.2/24 and 10.10.3.1/24. Then, configure area 0.0.0.0, which includes network 10.10.1.2/24 and stub area 0.0.0.3, which includes network 10.10.3.1/24.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.3.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 2.2.2.2
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-2 vif vlan-2 address
10.10.1.2
admin@XorPlus# set protocols ospf4 area 0.0.0.2 interface vlan-3 vif vlan-3 address
10.10.3.1
```

```

admin@XorPlus# set protocols ospf4 area 0.0.0.2 area-type stub
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure just one OSPF interface in area 0.0.0.2.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.9.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 0.0.0.2 interface vlan-2 vif vlan-2 address
10.10.3.2
admin@XorPlus# set protocols ospf4 area 0.0.0.2 area-type stub
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch D

Configure just one OSPF interface in area 0.0.0.1. Switch D should import the RIP or BGP route from the RIB and distribute it to other areas.

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.2.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.8.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set policy policy-statement rip-ospf term rip from protocol rip
admin@XorPlus# set policy policy-statement rip-ospf term rip then external-type 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4

```

```

admin@XorPlus# set protocols ospf4 area 0.0.0.1 interface vlan-2 vif vlan-2 address
10.10.2.2
admin@XorPlus# set protocols ospf4 area 0.0.0.1 area-type nssa
admin@XorPlus# set protocols ospf4 export rip-ospf
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

OSPF Stub Area NSSA Summary

By default, external routes and inter-area routes will be injected into stub areas or NSSAs. User can utilize the **summaries disable true** parameter to prevent external printer-area routes from being injected into stub areas or NSSAs. User can also use **set protocols ospf4 area <area-id> default-lsa disable false** to create a default route entry.

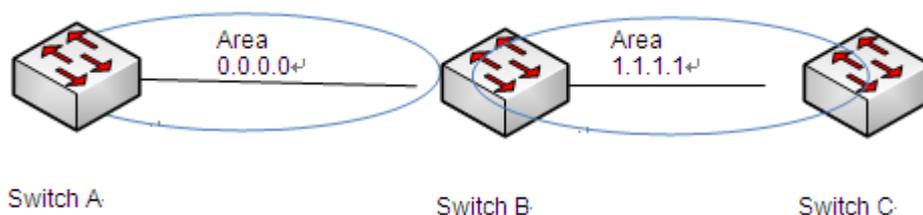


Figure 5-5. OSPF Stub area/NSSA summary: area 1.1.1.1 should be a stub area or an NSSA

Configuring Switch A

```

admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.2
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

OSPF Virtual Link Configuration Guide

Configure Virtual Links

The single backbone area (area 0.0.0.0), cannot be disconnected, or certain areas of the Autonomous System will become unreachable. To establish and maintain connectivity of the backbone, virtual links can be configured through non-backbone areas. Virtual links serve to connect physically separate components of the backbone.

The two endpoints of a virtual link are Area Border Routers (ARBs). The virtual link must be configured in both routers. The configuration information in each router consists of the other virtual endpoint (the other ARB), and the non-backbone area that the two routers have in common (called the transit area). Virtual links cannot be configured through stub areas.

Enable OSPF on Switch A, B, C, and D at the beginning. There is no route entry from the backbone area (0.0.0.0) to area 2.2.2.2.

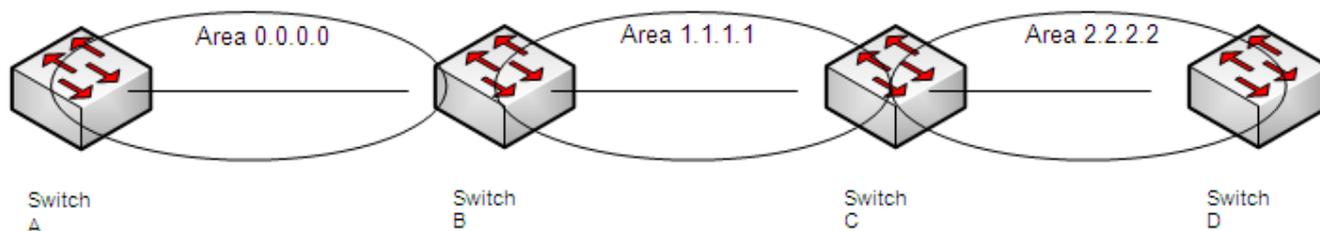


Figure 5-7. Virtual link configuration.

Configuring Switch A

```
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.2
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/47 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.2
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch B

```
admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.250
prefix-length 30
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.1
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/48 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.1
admin@XorPlus# set protocols ospf4 area 1.1.1.1 interface vlan-400 vif vlan-400 address
172.25.150.250
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch C

```

admin@XorPlus# set vlans vlan-id 300 l3-interface vlan-300
admin@XorPlus# set vlan-interface interface vlan-300 vif vlan-300 address 172.25.150.246
prefix-length 30
admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.249
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set interface gigabit-ethernet te-1/1/51 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 2.2.2.2 interface vlan-300 vif vlan-300 address
172.25.150.246
admin@XorPlus# set protocols ospf4 area 1.1.1.1 interface vlan-400 vif vlan-400 address
172.25.150.249
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Configuring Switch D

```

admin@XorPlus# set vlans vlan-id 300 l3-interface vlan-300
admin@XorPlus# set vlan-interface interface vlan-300 vif vlan-300 address 172.25.150.245
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set protocols ospf4 router-id 2.2.2.2
admin@XorPlus# set protocols ospf4 area 2.2.2.2 interface vlan-300 vif vlan-300 address
172.25.150.245
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Enable Virtual Links on the Area Border Routers

(Switch B and Switch C). After this step, there will be a route entry from the backbone area, **0.0.0.0**, to area **2.2.2.2**.

Configuring Switch B

```

admin@XorPlus# set protocols ospf6 area 0.0.0.0 virtual-link 3.3.3.3 transmit-area 1.1.1.1
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Configuring Switch C

```

admin@XorPlus# set protocols ospf6 area 0.0.0.0 virtual-link 4.4.4.4 transmit-area 1.1.1.1
admin@XorPlus#commit

```

```

Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Checking an IPv6 OSPF

Check ipv6 ospf neighbor on Switch B.

```

admin@XorPlus# run show ospf4 neighbor
Address Interface State Router ID Pri Dead
-----
192.168.1.2 vlan-500/vlan-500 Full 1.1.1.1 128 34
172.25.150.249 vlan-400/vlan-400 Full 3.3.3.3 128 36
172.25.150.249 vlink/3.3.3.3 Init 3.3.3.3 0 0

```

Configuring Switch B

```

admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.1
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-100 vif vlan-500 address
192.168.1.1
admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.250
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-400 vif vlan-400 address
172.25.150.250
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Configuring Switch C

```

admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.249
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-400 vif vlan-400 address
172.25.150.249
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Configuring Area 1.1.1.1 as a Stub area or NSSA

```
admin@XorPlus# set protocols ospf4 area 1.1.1.1 area-type <normal | stub | nssa>
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Check route table on DUT3, there will be route entry to backbone area 192.168.1.0/30

```
admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
172.25.150.248 255.255.255.252 08:9E:01:62:D5:61 connected
192.168.1.0 255.255.255. 252 60:EB:69:9B:BE:31 te-1/1/51
Total route count:2
```

Disabling the summary function on ABR(DUT2 area 1.1.1.1)

```
admin@XorPlus# set protocols ospf4 area 1.1.1.1 summaries disable true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Check route table on DUT3, the route entry to backbone area was lost

```
admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
172.25.150.248 255.255.255.252 08:9E:01:62:D5:61 connected
Total route count:1
```

Enable default-lsa function on ABR(DUT2)

```
admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
172.25.150.248 255.255.255.252 08:9E:01:62:D5:61 connected
0.0.0.0 0.0.0.0 60:EB:69:9B:BE:31 te-1/1/51
Total route count:2
```

OSPF Area Range Configuration Guide

OSPF should aggregate route entries from the backbone area into a non-backbone area or from a non-backbone area into the backbone area. Route aggregation works only on the ABR.

User can use the "advertise disable" parameter to restrain route aggregation at the ABR. The ABR will aggregate routes by default after **area-range** is configured and the packet is routed to the best (the longest or most specific) match.

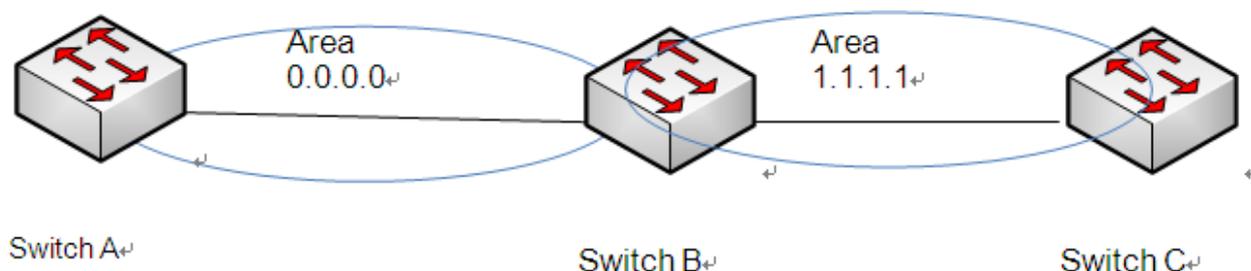


Figure 5-6. OSPF area range configuration.

Configuring Switch A

```
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.2
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.2
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch B

```
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.1
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.1
admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.250
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-400 vif vlan-400 address
172.25.150.250
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch C

```

admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.249
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-400 vif vlan-400 address
172.25.150.249
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Checking the route table on Switch C

There will be a 30-bit route entry, 192.168.1.0/30.

```

admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
172.25.150.248 255.255.255.252 08:9E:01:62:D5:61 connected
192.168.1.0 255.255.255. 252 60:EB:69:9B:BE:31 te-1/1/51
Total route count:2

```

Configuring area-range on ABR(DUT2)

```

admin@XorPlus# set protocols ospf4 area 0.0.0.0 area-range 192.168.1.0/24 advertise true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Checking the route table on DUT3

The route entry 192.168.1.0/30 will be replaced by 192.168.1.0/24.

```

admin@XorPlus# run show route forward-route ipv4 all
Destination NetMask NextHopMac Port
-----
172.25.150.248 255.255.255.252 08:9E:01:62:D5:61 connected
192.168.1.0 255.255.255. 0 60:EB:69:9B:BE:31 te-1/1/51
Total route count:2

```

Importing an External Route into an OSPF Area

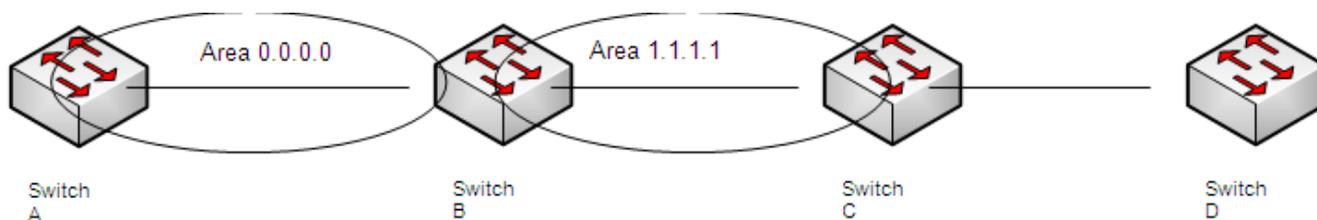


Figure 5-8. Importing an external route into an OSPF area.

Configuring Switch A

```
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.2
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/47 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.2
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch B

```
admin@XorPlus# set vlans vlan-id 400 l3-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.250
prefix-length 30
admin@XorPlus# set vlans vlan-id 500 l3-interface vlan-500
admin@XorPlus# set vlan-interface interface vlan-500 vif vlan-500 address 192.168.1.1
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/48 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 4.4.4.4
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-500 vif vlan-500 address
192.168.1.1
admin@XorPlus# set protocols ospf4 area 1.1.1.1 interface vlan-400 vif vlan-400 address
172.25.150.250
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#
```

Configuring Switch C

```
admin@XorPlus# set vlans vlan-id 300 l3-interface vlan-300
admin@XorPlus# set vlan-interface interface vlan-300 vif vlan-300 address 172.25.150.246
```

```

prefix-length 30
admin@XorPlus# set vlans vlan-id 400 13-interface vlan-400
admin@XorPlus# set vlan-interface interface vlan-400 vif vlan-400 address 172.25.150.249
prefix-length 30
admin@XorPlus# set interface gigabit-ethernet te-1/1/49 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set interface gigabit-ethernet te-1/1/51 family ethernet-switching
native-vlan-id 400
admin@XorPlus# set protocols ospf4 router-id 3.3.3.3
admin@XorPlus# set protocols ospf4 area 2.2.2.2 interface vlan-300 vif vlan-300 address
172.25.150.246
admin@XorPlus# set protocols ospf4 area 1.1.1.1 interface vlan-400 vif vlan-400 address
172.25.150.249
admin@XorPlus#commit
Waiting for merging configuration.
Commit OK.
Save Done.
admin@XorPlus#

```

Configuring an external route import policy on Switch C

1. Configure external static route.

```

admin@XorPlus# set protocols static route 192.168.6.0/24 next-hop 172.25.150.245
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

2. Configure policy to import external route

```

admin@XorPlus# set policy policy-statement static term 1 from protocol static
admin@XorPlus# set policy policy-statement static then accept
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

3. Use policy on ospf

```

admin@XorPlus# set protocols ospf4 export static
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

4. Check route table on Switch A. There will be route entry 192.168.6.0/24

```

admin@XorPlus# run show route forward-route ipv4 all

```

Destination	NetMask	NextHopMac	Port
192.168.1.0	255.255.255.252	C8:0A:A9:AE:0A:66	connected
172.25.150.248	255.255.255.252	60:EB:69:9B:BE:31	te-1/1/47
192.168.6.0	255.255.255.0	60:EB:69:9B:BE:31	te-1/1/47

```

Total route count:3

```

BFD Supports for OSPF, BGP, Static Route and ECMP

Configuring the Mode

There are two BFD modes: active and passive. The BFD will send protocol messages 'actively' in active mode and 'passively' in passive mode.

```
XorPlus# set protocols bfd mode active
XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
XorPlus# set protocols bfd mode passive
XorPlus# commit
Commit OK.
Save done.
XorPlus#
```

Configuring Detect-Multiplier, Min-Receive-Interval and Min-Transmit-Interval

Detect-multiplier: a detection timeout multiple. It is used in calculating detection timeout time by the detector.

Min-receive-interval: the minimum sending interval of the BFD packet supported by the local side.

Min-transmit-interval: the minimum receiving interval of the BFD packet supported by the local side.

```
XorPlus# set protocols bfd interface vlan25 detect-multiplier 5
XorPlus# set protocols bfd interface vlan25 min-transmit-interval 1000
XorPlus# set protocols bfd interface vlan25 min-receive-interval 2000
XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
XorPlus#
```

Enable BFD on L3 Interface

Enable BFD on the VLAN interface.

```
XorPlus# set protocols bfd interface vlan25 disable false
XorPlus# commit
Commit OK.
Save done.
XorPlus#
```

Enable BFD Supporting for OSPF4

Enable BFD to support for protocol OSPF4.

```
XorPlus# set protocols ospf4 area 1.1.1.1 interface vlan-25 vif vlan-25 address
125.125.25.6 bfd disable false
XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
XorPlus#
```

Enable BFD Supporting for OSPF6

Enable BFD to support for protocol OSPF6.

```
XorPlus# set protocols ospf6 area 1.1.1.1 interface vlan-23 vif vlan-23 bfd disable false
XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
XorPlus#
```

Enable BFD Supporting for BGP

Enable BFD to support for protocol BGP.

```
XorPlus# set protocols bgp peer 125.125.25.1 bfd disable false
XorPlus# commit
Merging the configuration.
Commit OK.
XorPlus#
```

Enable BFD Supporting for Static Route

Enable BFD to support for protocol static route.

```
XorPlus# set protocols static route 201.201.20.0/24 next-hop 113.113.13.1
XorPlus# set protocols static route 201.201.20.0/24 bfd true
XorPlus# commit
Commit OK.
Save done.
XorPlus#
```

Enable BFD Supporting for ECMP

Enable BFD to support for protocol ECMP.

```
XorPlus# set protocols static route 201.201.20.0/24 next-hop 113.113.13.1
XorPlus# set protocols static route 201.201.20.0/24 bfd true
XorPlus# commit
Commit OK.
Save done.
XorPlus# set protocols static route 201.201.20.0/24 qualified-next-hop 115.115.15.1 bfd
true
XorPlus# set protocols static route 201.201.20.0/24 qualified-next-hop 115.115.15.1 metric
1
XorPlus# commit
Commit OK.
Save done.
XorPlus#
```

BFD Configuration

Fig.5-9 presents an example of configuring BFD supporting for OSPF4. Switch A and Switch B are located in the backbone area 0.0.0.0.

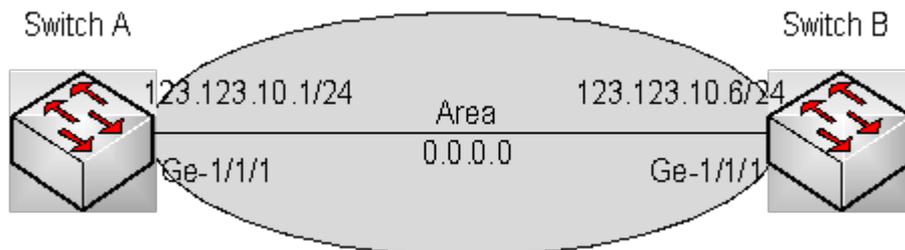


Figure 5-9. BFD basic configuration.

Configuring Switch A

For switch A, configure one VLAN interface for networks 123.123.10.1/24. Also configure area 0.0.0.0, which includes networks 123.123.10.1/24 and 123.123.10.6/24, and enable BFD on OSPF4 and VLAN interface.

```
admin@XorPlus# set vlans vlan-id 10
admin@XorPlus# set vlans vlan-id 10 13-interface vlan10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set vlan-interface interface vlan10 vif vlan10 address 123.123.10.1
prefix-length 24
```

```

admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 1.1.1.1
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan10 vif vlan10 address
123.123.10.1 bfd disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bfd interface vlan10 disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B

For switch B, configure one VLAN interface for networks 123.123.10.6/24. Also configure area 0.0.0.0, which includes networks 123.123.10.1/24 and 123.123.10.6/24, and enable BFD on OSPF4 and VLAN interface.

```

admin@XorPlus# set vlans vlan-id 10
admin@XorPlus# set vlans vlan-id 10 13-interface vlan-10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set vlan-interface interface vlan-10 vif vlan-10 address 123.123.10.6
prefix-length 24
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols ospf4 router-id 2.2.2.2
admin@XorPlus# set protocols ospf4 area 0.0.0.0 interface vlan-10 vif vlan-10 address
123.123.10.6 bfd disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bfd interface vlan-10 disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verifying the BFD Configuration

Verify the BFD configuration of a switch by checking its BFD neighbor.

```

admin@XorPlus# run show bfd neighbor ipv4
Detect Transmit
Local Address Remote Address Interface State Time(ms) Interval(ms) Multiplier
-----

```

```
123.123.10.1 123.123.10.6 vlan10 Up 1500 500 3
admin@XorPlus#
```

Fig.5-10 presents an example of configuring BFD supporting for static route.

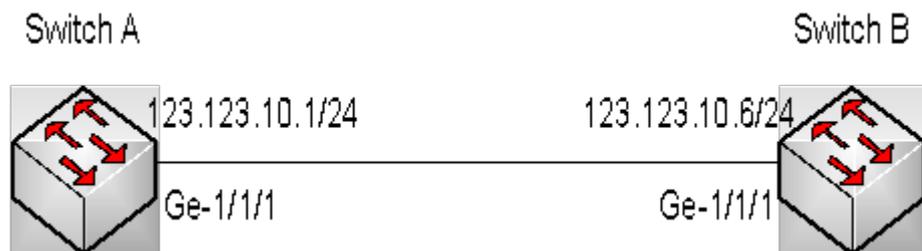


Figure 5-10. BFD basicconfiguration

Configuring Switch A

For switch A, configure one VLAN interface for networks 123.123.10.1/24. Also configure static route whose next hop direct to network 123.123.10.6/24, and enable BFD on static route and VLAN interface.

```
admin@XorPlus# set vlans vlan-id 10
admin@XorPlus# set vlans vlan-id 10 l3-interface vlan10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set vlan-interface interface vlan10 vif vlan10 address 123.123.10.1
prefix-length 24
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 200.200.10.0/24 next-hop 123.123.10.6
admin@XorPlus# set protocols static route 200.200.10.0/24 bfd true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bfd interface vlan10 disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B

For switch B, configure one VLAN interface for networks 123.123.10.6/24. Also configure static route whose next hop direct to network 123.123.10.1/24, and enable BFD on static route and VLAN interface.

```
admin@XorPlus# set vlans vlan-id 10
admin@XorPlus# set vlans vlan-id 10 l3-interface vlan-10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
```

```

admin@XorPlus# set vlan-interface interface vlan-10 vif vlan-10 address 123.123.10.6
prefix-length 24
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 178.178.10.0/24 next-hop 123.123.10.1
admin@XorPlus# set protocols static route 178.178.10.0/24 bfd true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bfd interface vlan-10 disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Verifying the BFD Configuration

Verify the BFD configuration of a switch by checking its BFD neighbor.

```

admin@XorPlus# run show bfd neighbor ipv4
Detect Transmit
Local Address Remote Address Interface State Time(ms) Interval(ms) Multiplier
-----
123.123.10.1 123.123.10.6 vlan10 Up 1500 500 3
admin@XorPlus#

```

Configuring ECMP (Equal-Cost Multipath Routing)

In L2/L3, ECMP is supported. The maximum ECMP outgoing port group is 4×128 . If each ECMP route is configured to have up to 4 equal-cost paths, for example, the maximum ECMP outgoing port group support is 128. If each ECMP route is configured to have up to 16 equal-cost paths, the maximum ECMP outgoing port group support is 32. Several different ECMP routes can share the same outgoing port group. After configuring the ECMP equal-cost path maximum, reboot the switch to make it available.

Configuring the Equal-Cost Path Maximum

```

admin@XorPlus# set interface ecmp path_max 8
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
ECMP max path changes, please reset the box!
admin@XorPlus# run request system reboot
The system is going down NOW!
Sending SIGTERM to all processes
Sending SIGKILL to all processes
Requesting system reboot
Restarting system.
rstcr compatible register does not exist!
uses the mpc8541's gpio to do a reset.

```

```

U-Boot 1.3.0 (Sep 8 2010 - 17:20:00)
CPU: 8541, Version: 1.1, (0x80720011)
Core: E500, Version: 2.0, (0x80200020)
Clock Configuration:
CPU: 825 MHz, CCB: 330 MHz,
DDR: 165 MHz, LBC: 41 MHz
L1: D-cache 32 kB enabled
I-cache 32 kB enabled
I2C: ready
DRAM: Initializing

```

Configuring Static ECMP Routing

```

admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set vlans vlan-id 4
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 4
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlans vlan-id 4 l3-interface vlan-4
admin@XorPlus# set vlan-interface interface vlan-2 address 10.10.60.10 prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 address 10.10.61.10 prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-4 address 10.10.62.10 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 10.10.51.0/24 next-hop 10.10.61.20
admin@XorPlus# set protocols static route 10.10.51.0/24 qualified-next-hop 10.10.62.20
metric 1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
You can che

```

Check the static ECMP route for 10.10.51.0/24 in the RIB.

```

admin@XorPlus# run show route table ipv4 unicast final
10.10.51.0/24 [static(1)/1]
> to 10.10.61.20 via vlan-3/vlan-3
10.10.51.0/24 [static(1)/1]
> to 10.10.62.20 via vlan-4/vlan-4
10.10.60.0/24 [connected(0)/0]
> via vlan-2/vlan-2
10.10.61.0/24 [connected(0)/0]
> via vlan-3/vlan-3
10.10.62.0/24 [connected(0)/0]
> via vlan-4/vlan-4

```

Configuring ECMP Hash Fields

In the default setting, all fields are hashed by IP-source, port-destination, port-source, and VLAN. Additional fields can be enabled, as shown below:

```
admin@XorPlus# set interface ecmp hash-mapping field ingress-interface disable false
admin@XorPlus# set interface ecmp hash-mapping field ip-destination disable false
admin@XorPlus# set interface ecmp hash-mapping field ip-protocol disable false
admin@XorPlus# set interface ecmp hash-mapping field ip-source disable false
admin@XorPlus# set interface ecmp hash-mapping field port-destination disable false
admin@XorPlus# set interface ecmp hash-mapping field port-source disable false
admin@XorPlus# set interface ecmp hash-mapping field vlan disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring VRRP (Virtual Router Redundancy Protocol)

Configuring VRRP

In L2/L3, VRRP is supported for both preempt and non-preempt parameters.

In the configuration below, a virtual router with IP 192.168.1.5/24 is created. VRRP preemption and VRRP priority can be configured.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols vrrp interface vlan-2 vrid 1
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 ip 192.168.1.5
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 preempt true
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 priority 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Check the VRRP configuration.

```
admin@XorPlus# run show vrrp vlan-2
Interface vlan-2
```

```
Vif vlan-2
VRID 1
State master
Master IP 192.168.1.1
admin@XorPlus#
```

Configuring Active-Active-VRRP

In L2/L3, Active-Active-VRRP is supported for both preempt and non-preempt parameters.

The only different between Vrrp and Active-Active-VRRP command is that Active-Active-VRRP should enable load-balance.

Active-Active-VRRP can produce two virtual mac addresses, the format of master virtual mac address is 00:00:5E:00:01:VRID, the format of slave virtual mac address is 00:00:5E:00:02:VRID. These two virtual mac should be installed on both master and slave switches. These virtual mac is used for communicating with each other(master and slave) and other equipment. When the host send arp request for virtual ip, the master node will send arp reply with one of two virtual mac addresses by hash method based on the source mac address in the arp request packet.

When the arp table learned a new mac address of the host on the one of two vrrp nodes, this node will notice the another node with the new ip address, and then the another node will send arp request for the mac address of the host. The max 255 ip addresses are announced at a time between master and slave node.

How to deal with arp request for Virtual Ip

The master can not learn the arp when the arp request destination ip address is virtual ip. But it can send arp reply for this arp request.

The rules of synchronous Arp

Whether the Arp table entry can be synchronized is decided by the ip in arp table entry. If the ip network segment are same with virtual ip network segment.

Suggestion

It is better to configure the virtual ip network segment same with the L3 interface ip network segment.

Notice

Active-Active-VRRP only can support two VRRP node, one is master, another is slave.

In the configuration below, a virtual router with IP 192.168.1.5/24 is created. VRRP preemption and VRRP priority can be configured.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 192.168.1.1
prefix-length 24
```

```

admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 192.168.2.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols vrrp interface vlan-2 vrid 1
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 ip 192.168.1.5
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 load-balance disable
false
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 preempt true
admin@XorPlus# set protocols vrrp interface vlan-2 vif vlan-2 vrid 1 priority 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Check the VRRP configuration.

```

admin@XorPlus# run show vrrp vlan-2
Interface vlan-2
Vif vlan-2
VRID 1
State master
Master IP 192.168.1.1
admin@XorPlus#

```

Configuring the IPv6 Neighbor Aging Time

Configure the IPv6 neighbor aging time. The neighbor will be removed after the timer has expired.

```

admin@XorPlus# set protocols neighbour aging-time 480
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring a Static IPv6 Neighbor

Configure a static IPv6 neighbor in a specified interface.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set protocols neighbour interface vlan-2 vif vlan-2 address 2001::01
mac-address 22:22:22:22:22:22
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show ipv6-neighbors static
aging-time(seconds): 480
Address HW Address Interface
-----
2001::1 22:22:22:22:22:22 vlan-2
admin@XorPlus#

```

Configuring IPv6 Router A dvertisement

Enable router advertisement messages.

```
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlan-interface interface vlan1 router-advertisement disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

IPv6 Static Routing Configuration

In L2/L3, IPv6 static routing is supported. The IPv6 for OSPFv3 and RIPng will be supported soon.

In P-3290 and P-3780, configure the link-local IPv6 address, otherwise all the IPv6 interfaces will share the same link-local address. This problem will be fixed in future versions.

Configuring a Static Route for IPv6

Configure the link-local address and global address for a VLAN interface.

```
admin@XorPlus# set vlans vlan-id 2
admin@XorPlus# set vlans vlan-id 3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address
2001:db8:3c4d:5:60:ff:73:87 prefix-length 64
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address
fe80::ca0a:a9ff:fe04:4931 prefix-length 64
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address
2001:db8:3c4d:6:0:ff:73:87 prefix-length 64
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address fe80::ca0a:a9ff:4:4932
prefix-length 64
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 2001:db8:3c4d:7::/64 next-hop
2001:db8:3c4d:5:60:d6ff:73:89
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Then, verify the IPv6 static route in the RIB:

```

admin@XorPlus# run show route table ipv6 unicast final
2001:db8:3c4d:5::/64[connected(0)/0]
> via vlan-2/vlan-2
2001:db8:3c4d:6::/64[connected(0)/0]
> via vlan-3/vlan-3
fe80::/64 [connected(0)/0]
> via vlan-3/vlan-3
fe80::/64 [connected(0)/0]
> via vlan-2/vlan-2

```

OSPFv3 Routing Protocol Configuration

In admin@XorPlus, OSPFv3 is supported.

Configuring the Router ID

```

admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configuring an OSPF Area and Area-Type

Area 0.0.0.0 is the backbone area of OSPF; each OSPF domain should have the area 0.0.0.0. Area types includes normal, stub, and NSSA.

```

admin@XorPlus# set protocols ospf6 area 0.0.0.0 area-type normal
admin@XorPlus# set protocols ospf6 area 1.1.1.1 area-type stub
admin@XorPlus# set protocols ospf6 area 2.2.2.2 area-type nssa
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring OSPF Interfaces

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 2001::15
prefix-length 64
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 2002::15
prefix-length 64
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan-2 vif vlan-2 address

```

```

2001::15
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan-3 vif vlan-3 address
2002::15
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 1.1.1.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show ospf6 interface
Interface State Area DR ID BDR ID Nbrs
-----
vlan-2 Down 0.0.0.0 0.0.0.0 0.0.0.0 0
vlan-3 Down 0.0.0.0 0.0.0.0 0.0.0.0 0

```

Configuring Additional OSPF Interface Parameters

See below to configure additional OSPF interface parameters (hello-interval, interface-cost, static neighbor, priority, retransmit-interval, router-dead-interval, and transmit-delay).

```

admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan-2 vif vlan-2 hello-interval
10
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan-2 vif vlan-2 interface-cost
8
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan-2 vif vlan-2 transmit-delay
2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

ACL and Filter Configuration

In L2/L3, ACLs support destination-address-ipv4, destination-address-ipv6, destination-mac-address, destination-port, ether-type, ip, protocol, source-address-ipv4, source-address-ipv6, source-mac-address, source-port, and vlan-id.

TCP flags are also supported. These ACLs can be applied to physical ports, LAG ports, and VLAN interfaces. One ACL can be applied to multiple ports (the properties of the ports can be same or different), but only one port can be matched to one ACL.



ACL can't filter layer 2 protocol packets, for example BPDU, LLDP, LACP and so on.

Configuring ACLs

```

admin@XorPlus# set firewall filter bad-net sequence bad-1 from source-address-ipv4
1.1.1.0/24
admin@XorPlus# set firewall filter bad-net sequence bad-1 then action discard
admin@XorPlus# set firewall filter bad-net sequence bad-2 from source-address-ipv4
1.1.2.0/24
admin@XorPlus# set firewall filter bad-net sequence bad-2 then action discard

```

```

admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set firewall filter bad-net input interface ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set firewall filter bad-net input interface ael
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

When the switch receives a packet in ingress and egress, it will attempt to match ACLs by sequence number, with smaller values representing higher priorities. If the matched ACL's action is "forward" or "discard," the switch will forward or discard the packet and will not match the remaining ACLs. If there is no matching ACL, the packet will be dropped.

Configuring ACLs in VLANs

Every member port in the VLAN interface will be applied with the ACLs configured in the VLAN interface.

```

admin@XorPlus# set firewall filter bad-net sequence bad-1 from source-address-ipv4
1.1.1.0/24
admin@XorPlus# set firewall filter bad-net sequence bad-1 then action discard
admin@XorPlus# set firewall filter bad-net sequence bad-2 from source-address-ipv4
1.1.2.0/24
admin@XorPlus# set firewall filter bad-net sequence bad-2 then action discard
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set firewall filter bad-netinput vlan-interface vlan-2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring ACL Discard TCP ACK

You can configure ACL TCP flags (ACK/FIN/PSH/RST/SYN/URG/TCP-ESTABLISHED/TCP-INITIAL) to specify what action (forward/discard) to perform on which packets (true/false).

```

admin@XorPlus# set firewall filter bad-net sequence bad-1 then action discard
admin@XorPlus# set firewall filter bad-net sequence bad-1 from protocol tcp flags ack true
admin@XorPlus# commit
Waiting for merging configuration.

```

```

Commit OK.
Save done.
admin@XorPlus# set firewall filter bad-net output interface ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring ACL logging for Match Statistics

```

admin@XorPlus# set firewall filter bad-net sequence bad-1 then action discard
admin@XorPlus# set firewall filter bad-net sequence bad-1 from destination-address-ipv4
192.168.100.0/24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set firewall filter bad-net input interface ge-1/1/1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set firewall filter bad-net sequence bad-1 log interval 10
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run syslog monitor on
admin@XorPlus#

```

- [Configuring Control Plane Security Guide](#)
 - [Configuration Control Plane Security Policer](#)
 - [Control Plane Security Examples](#)

Configuring Control Plane Security Guide

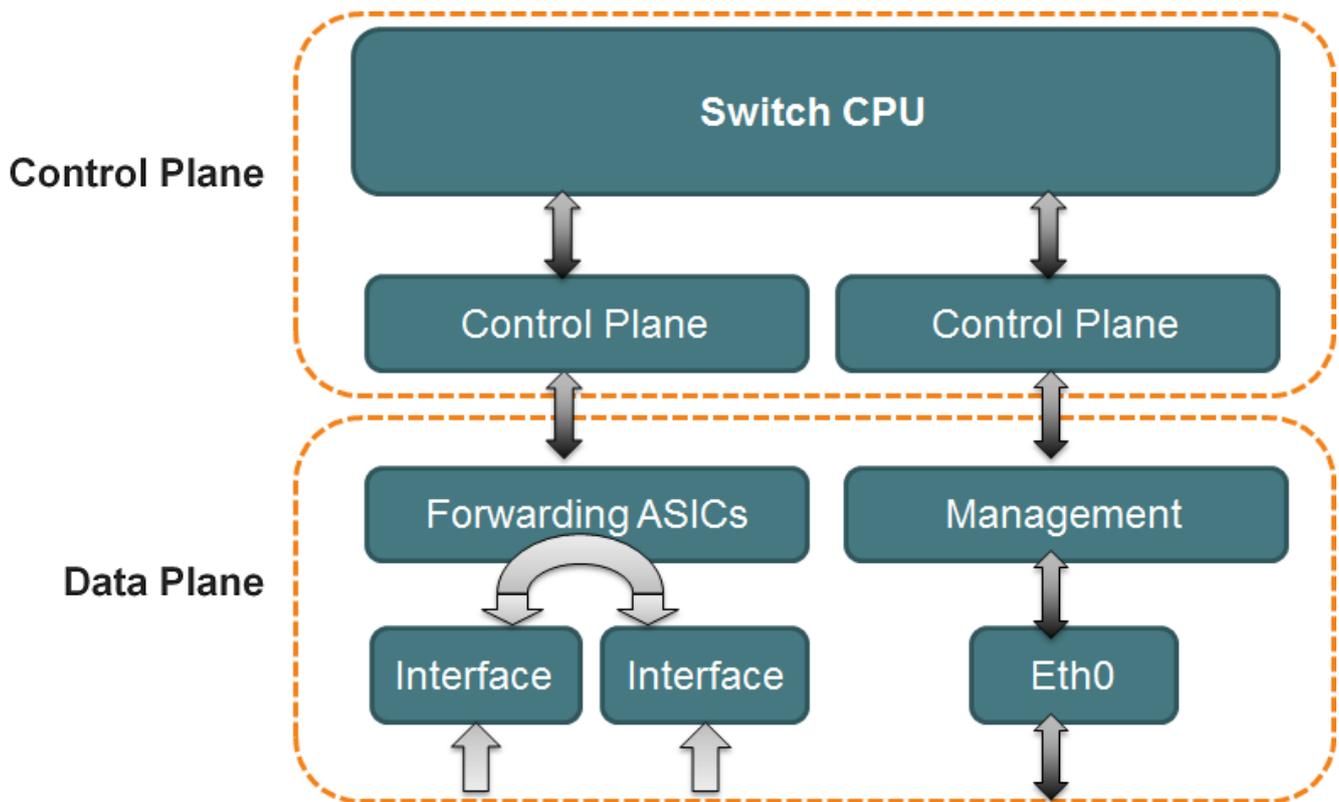
- [Configuration Control Plane Security Policer](#)
- [Control Plane Security Examples](#)

About Control Plane Security Policer

Control plane security is an firewall policer application on PicOS switches that allows the configuration of firewall policies that rate-limit the traffic handled by the main CPU of the network device. This protects the control plane of the switch from direct denial-of-service attacks. With Control plane security, these firewall policies are configured to permit, block, or rate-limit the packets handled by the main CPU.

Control plane security can be applied in management interface and inbound interface. Create a firewall filter on the ASIC to protect the Control Plane and modify the IPFilter for inband interface, and create IPFilter configuration for management interface.

Figure illustrates the flow of packets from various interfaces. Packets destined to the control plane are subject to control plane filter checking, as depicted by the control plane services block.



Control Plane Security Policer Implementation

When configuring the control plane security policer, the sequence of events is as follows:

1. A packet enters the switch configured with control plane security on the ingress port.
2. The port performs any applicable input port and firewall filter services.
3. The packet gets forwarded to the switch CPU.
4. The switch CPU makes a routing or a switching decision, determining whether the packet is destined for the control plane.
5. Packets destined for the control plane are processed by control plane security and are dropped or delivered to the control plane according to each firewall filter policy. Packets that have other destinations are forwarded normally.

About Rate-limit and Burst applied in firewall policer

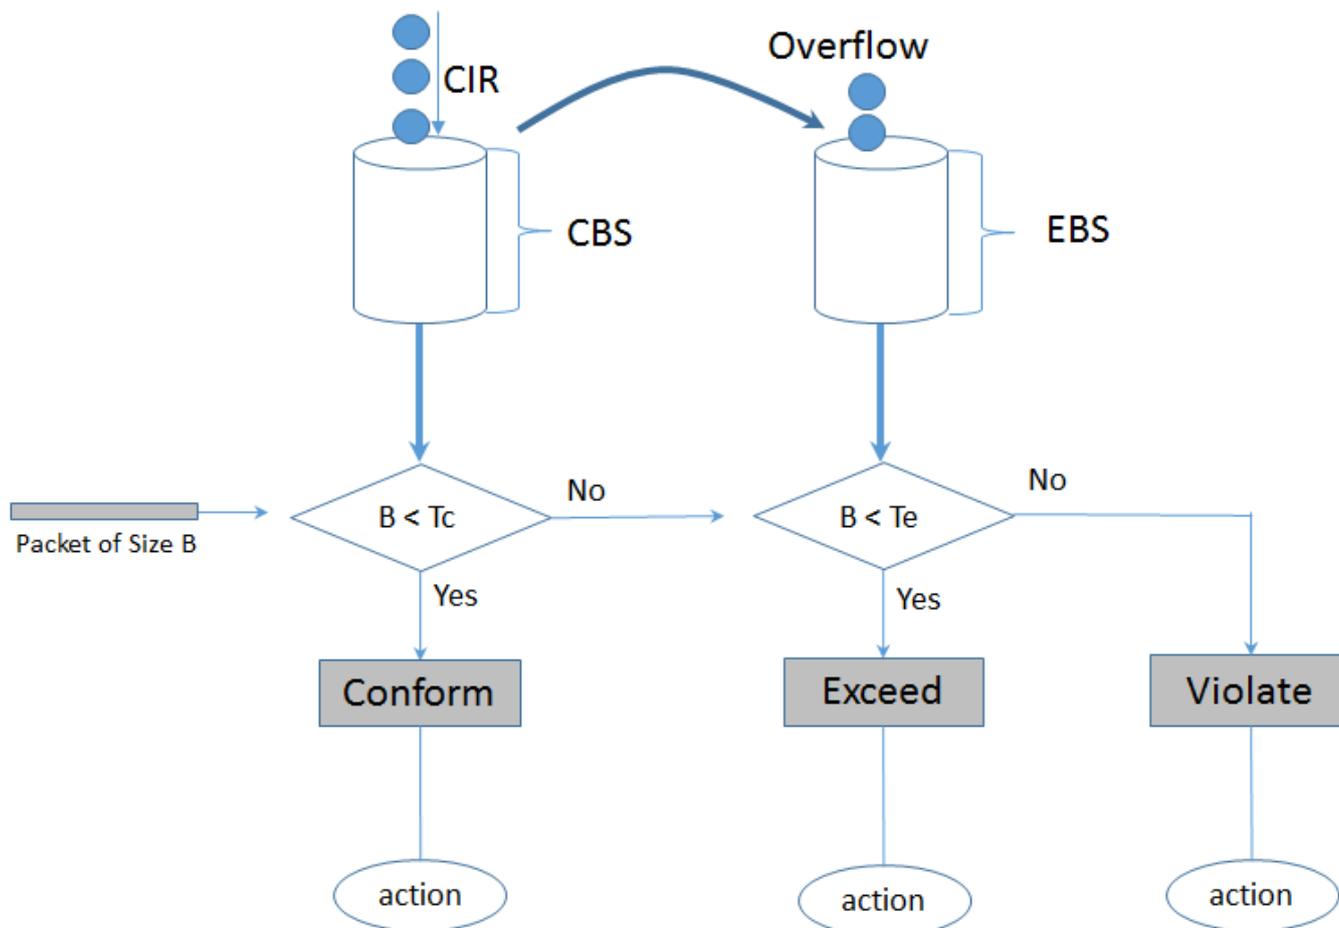


Figure. RFC2697 Single-Rate Three-Color Policer Logic

The single-rate three-color policer/marker algorithm with dual buckets is used to implement CoPP rate-limit and burst. Unlike the standard algorithm, the yellow traffic will be discarded as well as red traffic in order to make the implementation simple enough. Additionally, the action to be applied to a packet, forward or discard, totally depends on the T_c , the token counter of CBS, the instantaneous number of tokens left in the CBS bucket because both yellow and red traffic will be discarded for pica8 switch, which means the second bucket - EBS bucket - is not used to police the traffic at all. For example :

```
set firewall policer 10pps if-exceeding rate-limit 10 set firewall policer 10pps if-exceeding burst-limit 5
```

The above configuration is equal to below :

CIR (Committed Information Rate): 10 pps

CBS (Committed Burst Size): 5 packet

EBS (Excess Birst Size): 5 packet

10 pps is rate-limit value. 5 packet is burst value, the size of the bucket.

Control Plane Security Configuration Steps:

Step 1: Configure Control Plane Security policer rate-limit

```
set firewall policer 100pps if-exceeding rate-limit 100
set firewall policer 100pps if-exceeding burst-limit 5
```

```
set firewall policer 100pps then action discard
commit
```

Step 2: Configure firewall filter match and applied control plane security policer to firewall filter action

```
set firewall filter f1 sequence 1 from protocol icmp
set firewall filter f1 sequence 1 then policer 100pps
set firewall filter f1 sequence 1 then action forward
commit
```

Step 3: Configure firewall filter applied to inbound-control-plane or management interface (eth0,eth1)

```
set firewall filter f1 input interface inbound-control-plane
set firewall filter f1 input interface eth1
commit
```

Control Plane Security Rate Limits and Action Examples

The purpose of defining as follows:

- The BGP or IGP traffic is not to rate limit.
- The management protocol traffic is not to rate limit
- The arp, icmp traffic with rate limit

Filter class	Rate-limit(pps)	Conform Action	Exceed Action
ARP	100	Transmit	Drop
ICMP	100	Transmit	Drop
OSPF	null	Transmit	Transmit
BGP	null	Transmit	Transmit
ssh	null	Transmit	Transmit
telnet	null	Transmit	Transmit
Tacacs	null	Transmit	Transmit
ntp	null	Transmit	Transmit
snmp	null	Transmit	Transmit
default	200	Transmit	Drop

Configuring Control Plane Security Policer

```
set firewall policer 100pps if-exceeding rate-limit 100
set firewall policer 100pps if-exceeding burst-limit 5
set firewall policer 100pps then action discard
set firewall policer 200pps if-exceeding rate-limit 200
```

```

set firewall policer 200pps if-exceeding burst-limit 10
set firewall policer 200pps then action discard
commit

```

Configuring Control Plane Security Filter

```

set firewall filter f1 sequence 0 from ether-type 2054
set firewall filter f1 sequence 0 then policer 100pps
set firewall filter f1 sequence 0 then action forward
set firewall filter f1 sequence 1 from protocol icmp
set firewall filter f1 sequence 1 then policer 100pps
set firewall filter f1 sequence 1 then action forward
set firewall filter f1 sequence 2 from protocol ospf
set firewall filter f1 sequence 2 from destination-address-ipv4 224.0.0.5/32
set firewall filter f1 sequence 2 then action forward
set firewall filter f1 sequence 3 from protocol ospf
set firewall filter f1 sequence 3 from destination-address-ipv4 224.0.0.6/32
set firewall filter f1 sequence 3 then action forward
set firewall filter f1 sequence 4 from protocol tcp
set firewall filter f1 sequence 4 from source-port 179
set firewall filter f1 sequence 4 then action forward
set firewall filter f1 sequence 5 from protocol tcp
set firewall filter f1 sequence 5 from destination-port 179
set firewall filter f1 sequence 5 then action forward
set firewall filter f1 sequence 6 from protocol tcp
set firewall filter f1 sequence 6 from source-port 22
set firewall filter f1 sequence 6 then action forward
set firewall filter f1 sequence 7 from protocol tcp
set firewall filter f1 sequence 7 from destination-port 22
set firewall filter f1 sequence 7 then action forward
set firewall filter f1 sequence 8 from protocol tcp
set firewall filter f1 sequence 8 from source-port 23
set firewall filter f1 sequence 8 then action forward
set firewall filter f1 sequence 9 from protocol tcp
set firewall filter f1 sequence 9 from destination-port 23
set firewall filter f1 sequence 9 then action forward
set firewall filter f1 sequence 10 from protocol tcp
set firewall filter f1 sequence 10 from source-port 49
set firewall filter f1 sequence 10 then action forward
set firewall filter f1 sequence 11 from protocol tcp
set firewall filter f1 sequence 11 from destination-port 49
set firewall filter f1 sequence 11 then action forward
set firewall filter f1 sequence 12 from protocol udp
set firewall filter f1 sequence 12 from source-port 123
set firewall filter f1 sequence 12 then action forward
set firewall filter f1 sequence 13 from protocol udp
set firewall filter f1 sequence 13 from destination-port 123
set firewall filter f1 sequence 13 then action forward
set firewall filter f1 sequence 14 from protocol udp
set firewall filter f1 sequence 14 from source-port 161
set firewall filter f1 sequence 14 then action forward
set firewall filter f1 sequence 15 from protocol udp
set firewall filter f1 sequence 15 from destination-port 161
set firewall filter f1 sequence 15 then action forward
set firewall filter f1 sequence 16 from protocol udp
set firewall filter f1 sequence 16 from source-port 162
set firewall filter f1 sequence 16 then action forward
set firewall filter f1 sequence 17 from protocol udp
set firewall filter f1 sequence 17 from destination-port 162
set firewall filter f1 sequence 17 then action forward
set firewall filter f1 sequence 100 from
set firewall filter f1 sequence 100 then policer 200pps
set firewall filter f1 sequence 100 then action forward
commit

```

Configuring Control Plane Security Applied to Inbound-Interface

```
set firewall filter t1 input interface inbound-control-plane
commit
```

IPv4/IPv6 BGP Configuration

BGP protocol

- IPv4 BGP configuration
 - BGP Configuration Guide
 - BGP Basic Configuration Example
 - BGP Route Reflector Configuration Example
 - BGP Confederation Configuration Example
 - BGP Load Balancing Configuration Example
- IPv6 BGP Configuration
 - IPv6 BGP Introduction
 - Building Peering Sessions
 - EBGPeering
 - IBGP Peering
 - Establish BGP Peer Use 4-byte-AS-Number
 - Sources of Routing Updates
 - Injecting Information Dynamically into BGP
 - Injecting Information Statically into BGP
 - BGP Attributes
 - The NEXT_HOP Attribute
 - The AS_PATH Attribute
 - The LOCAL_PREF Attribute
 - The MULTI_EXIT_DISC Attribute
 - The COMMUNITY Attribute
 - BGP-4 Aggregation
 - Synchronization
 - Controlling Large-Scale Autonomous System
 - Confederations
 - Route Reflectors

- Redundancy and Load Balancing
- Designing Stable Internets
- Label BGP
 - Labeled BGP Support
 - Configuration Example for Labeled Support
 - IPV4 Labeled BGP Configuration
 - IPV6 Labeled-BGP Configuration
 - Debugging CLI for Labeled-BGP

IPv4 BGP configuration

IPv4 BGP configuration

- BGP Configuration Guide
- BGP Basic Configuration Example
- BGP Route Reflector Configuration Example
- BGP Confederation Configuration Example
- BGP Load Balancing Configuration Example

BGP Configuration Guide

Configuring a BGP Router ID

The router ID should be configured first when you configure BGP. The router ID is a string similar to the IP address, and is the identifier of a BGP router in an AS. You should not change the router ID after completing the configuration. By default, the BGP router ID is not configured.

```
admin@XorPlus# set protocols bgp bgp-id 1.1.1.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP Local-AS

The local AS (autonomous system) should be configured first when you configure BGP. The AS_Path attribute records to all the AS's that a route passes through, from the source to the destination, following the order of vectors.

```
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
```

```
Save done.  
admin@XorPlus#
```

Configuring External BGP Peering

If the AS number of the specified peer is different from the local AS number during the configuration of BGP peers, an EBGP peer is configured. To establish point-to-point connections between peer autonomous systems, configure a BGP session on each interface of a point-to-point link. Generally, such sessions are made at network exit points with neighboring hosts outside the AS.

```
admin@XorPlus# set protocols bgp local-as 100  
admin@XorPlus# set protocols bgp peer 192.168.49.1 as 200  
admin@XorPlus# set protocols bgp peer 192.168.49.1 next-hop-self true  
admin@XorPlus# set protocols bgp peer 192.168.49.1 local-ip 192.168.49.2  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring Internal BGP Peering

If the AS number of the specified peer is the same as the local AS number during the configuration of BGP peers, an IBGP peer is configured.

```
admin@XorPlus# set protocols bgp local-as 100  
admin@XorPlus# set protocols bgp peer 192.168.49.1 as 100  
admin@XorPlus# set protocols bgp peer 192.168.49.1 next-hop-self true  
admin@XorPlus# set protocols bgp peer 192.168.49.1 local-ip 192.168.49.2  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring the BGP Local Preference

Internal BGP (IBGP) sessions use a metric called the local preference, which is carried in IBGP update packets in the path attribute LOCAL_PREF. When an autonomous system (AS) has multiple routes to another AS, the local preference indicates the degree of preference for one route over the other routes. The route with the highest local preference value is preferred.

```
admin@XorPlus# set policy policy-statement send-network term t1 from network4  
172.168.200.0/24  
admin@XorPlus# set policy policy-statement send-network term t1 from protocol bgp  
admin@XorPlus# set policy policy-statement send-network term t1 then localpref 200  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.
```

```
admin@XorPlus# set protocols bgp peer 192.168.49.1 export send-network
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP MED

The multi-exit discriminator (MED) helps determine the optimal route for the incoming traffic of an AS, and is similar to the metric used in IGP. When a BGP device obtains multiple routes to the same destination address but with different next hops from EBGP peers, the BGP device selects the route with the smallest MED value as the optimal route.

```
admin@XorPlus# set policy policy-statement send-network term t1 from network4
172.168.200.0/24
admin@XorPlus# set policy policy-statement send-network term t1 from protocol bgp
admin@XorPlus# set policy policy-statement send-network term t1 then med 200
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bgp peer 192.168.49.1 export send-network
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP Next Hop

When an Autonomous System Boundary Router (ASBR) forwards the route learned from an EBGP peer to an IBGP peer, the ASBR, by default, does not change the next hop of the route. When the IBGP peer receives this route, it finds the next hop unreachable, sets the route to inactive, and does not use this route to guide traffic forwarding.

To enable the IBGP peer to use this route to guide traffic forwarding, configure the ASBR to set its IP address as the next hop of the route when the ASBR forwards this route to the IBGP peer. After the IBGP peer receives this route from the ASBR, it finds the next hop of the route reachable, sets the route to active, and uses this route to guide traffic forwarding.

When a BGP route changes, BGP needs to iterate the indirect next hop of the route again. If no restriction is imposed on the iterated route, BGP may iterate the next hop to an incorrect forwarding path, causing traffic loss. Configure routing policy-based route iteration to prevent traffic loss.

```
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 192.168.49.1 as 100
admin@XorPlus# set protocols bgp peer 192.168.49.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.49.1 local-ip 192.168.49.2
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring BGP Route Reflectors

To ensure the connectivity between IBGP peers within an AS, you need to establish full-mesh connections between the IBGP peers. When there are many IBGP peers, it is costly to establish a fully meshed network. A route reflector (RR) can solve this problem.

A cluster ID can help prevent routing loops between multiple RRs within a cluster, and between clusters. When a cluster has multiple RRs, the same cluster ID must be configured for all RRs within the cluster.

If full-mesh IBGP connections are established between clients of multiple RRs, route reflection between clients is not required and wastes bandwidth resources. In this case, prohibit route reflection between clients to reduce the network burden.

Within an AS, an RR transmits routing information and forwards traffic. When an RR connects to a large number of clients and non-clients, many CPU resources are consumed if the RR transmits routing information and forwards traffic simultaneously. This also reduces route transmission efficiency. To improve route transmission efficiency, prohibit BGP from adding preferred routes to IP routing tables on the RR, enabling the RR to only transmit routing information.

```
admin@XorPlus# set protocols bgp local-as 100  
admin@XorPlus# set protocols bgp peer 192.168.49.1 as 100  
admin@XorPlus# set protocols bgp peer 192.168.49.1 next-hop-self true  
admin@XorPlus# set protocols bgp peer 192.168.49.1 local-ip 192.168.49.2  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#  
admin@XorPlus# set protocols bgp route-reflector cluster-id 16.16.16.16  
admin@XorPlus# set protocols bgp route-reflector disable false  
admin@XorPlus# set protocols bgp peer 192.168.49.1 client true  
admin@XorPlus# commit  
Waiting for merging configuration.  
Commit OK.  
Save done.  
admin@XorPlus#
```

Configuring BGP Confederations

A confederation divides an AS into sub-AS's, which establish EBGP connections. Within each sub-AS, IBGP peers establish full-mesh connections or have an RR configured. On a large BGP network, configuring a confederation can reduce the number of IBGP connections, simplify routing policy management, and improve route advertisement efficiency.

```
admin@XorPlus# set protocols bgp local-as 65533  
admin@XorPlus# set protocols bgp peer 192.168.49.1 as 65533  
admin@XorPlus# set protocols bgp peer 192.168.49.1 next-hop-self true  
admin@XorPlus# set protocols bgp peer 192.168.49.1 local-ip 192.168.49.2  
admin@XorPlus# commit
```

```
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols bgp confederation identifier 2000
admin@XorPlus# set protocols bgp confederation disable false
admin@XorPlus# set protocols bgp peer 192.168.49.1 confederation-member true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring the BGP Connect Timer

Hold timers can be configured for all peers. The proper maximum interval at which 'keep alive' messages are sent is one third the hold time.

```
admin@XorPlus# set protocols bgp peer 192.168.49.1 holdtime 30
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring MD5 Authentication for TCP Connections

Configure Message Digest5 (MD5) authentication on a TCP connection between two BGP peers. The two peers must have the same configured password to establish TCP connections.

```
admin@XorPlus# set protocols bgp peer 192.168.11.10 md5-password pica8
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring EBGPFast-External-Fallover

This feature allows BGP to immediately respond to a fault on an interface, and delete the direct EBGPF sessions on the interface without waiting for the hold timer to expire. It implements rapid BGP network convergence. By default, EBGPFast-external-fallover is disabled.

```
admin@XorPlus# set protocols bgp fast-external-fallover disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP Route Summarization

BGP supports both automatic and manual route summarization. Manual route summarization takes precedence over automatic.

Configure automatic route summarization as follows:

```
admin@XorPlus# set protocols bgp auto-summary true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

BGP auto-summary true summarizes the routes exported by BGP.

To configure manual route summarization:

```
admin@XorPlus# set protocols bgp aggregate network4 192.168.1.0/24 suppress-detail true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP to Advertise Default Routes to Peers

The BGP device can be configured to send only a default route, with the local address as the next hop address, to its peer, regardless of whether there are default routes in the local routing table.

```
admin@XorPlus# set protocols bgp peer 192.168.11.10 default-route-advertise
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP to Remove Private AS Numbers

Private autonomous system (AS) numbers that range from 64512 to 65535 are used to conserve globally unique AS numbers. BGP can remove private AS numbers from updates to a peer.

```
admin@XorPlus# set protocols bgp peer 192.168.11.10 public-as-only
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP AS Loop

Repeated local AS numbers are allowed in routes. In the default setting, however, repeated local AS numbers are not allowed.

```
admin@XorPlus# set protocols bgp peer 192.168.11.10 allow-as-loop true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring BGP Load Balancing

If multiple paths to a destination exist, you can configure load balancing over such paths to improve link utilization.

Enable BGP load balancing:

```
admin@XorPlus# set protocols bgp multipath disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

BGP will not load balance across multiple paths by default. This is acceptable if you are multi-homed to a single AS, but if multi-homed to different AS path, you cannot load balance across theoretically equal paths.

Enter the **BGP multipath path-relax** command:

```
admin@XorPlus# set protocols bgp multipath path-relax true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuration Example

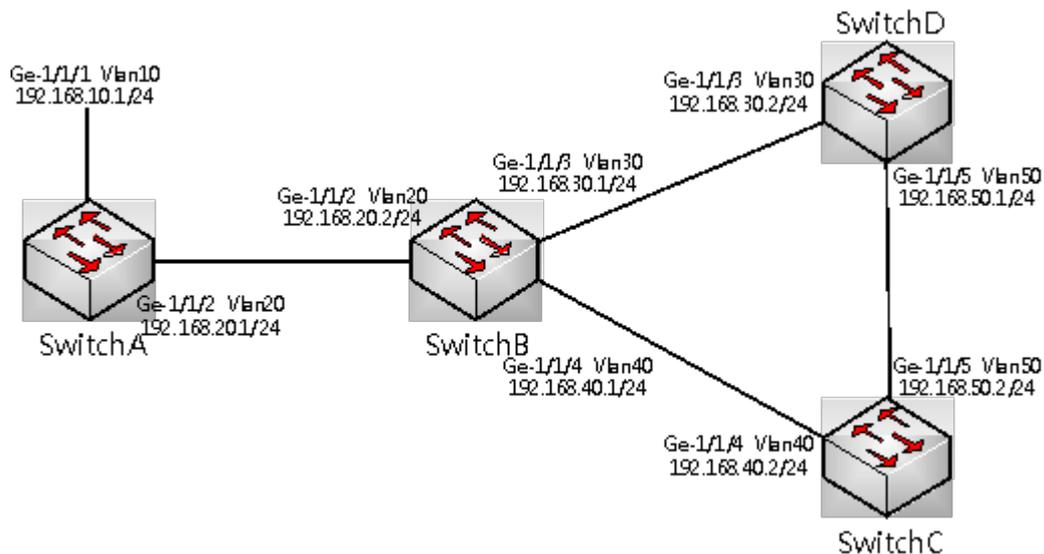
As shown in Fig. 5-9, BGP runs between switches.

An EBGP connection is established between Switch A and Switch B, and IBGP fullmesh connections are established between Switch B, Switch C, and Switch D.

Configure IBGP connections between Switch B, Switch C, and Switch D.

Configure an EBGP connection between Switch A and Switch B.

Figure 5-9. BGP configuration



Configuring Switch A

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 10 13-interface 10
admin@XorPlus# set vlans vlan-id 20 13-interface 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them IP addresses.

```
admin@XorPlus# set vlan-interface interface 10 vif 10 address 192.168.10.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP connection.

```
admin@XorPlus# set protocols bgp bgp-id 1.1.1.1
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.20.2 local-ip 192.168.20.1
```

```
admin@XorPlus# set protocols bgp peer 192.168.20.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 20 l3-interface 20
admin@XorPlus# set vlans vlan-id 30 l3-interface 30
admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces and assign them IP addresses.

```
admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 30 vif 30 address 192.168.30.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure EBGP and IBGP connections.

```
admin@XorPlus# set protocols bgp bgp-id 2.2.2.2
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.20.1 as 100
admin@XorPlus# set protocols bgp peer 192.168.20.1 local-ip 192.168.20.2
admin@XorPlus# set protocols bgp peer 192.168.20.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.30.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.30.2 local-ip 192.168.30.1
admin@XorPlus# set protocols bgp peer 192.168.30.2 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.40.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.40.2 local-ip 192.168.40.1
admin@XorPlus# set protocols bgp peer 192.168.40.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch C

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# set vlans vlan-id 50 l3-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them IP addresses.

```
admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.2 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an IBGP connection.

```
admin@XorPlus# set protocols bgp bgp-id 3.3.3.3
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.40.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.40.1 local-ip 192.168.40.2
admin@XorPlus# set protocols bgp peer 192.168.40.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.50.1 local-ip 192.168.50.2
admin@XorPlus# set protocols bgp peer 192.168.50.1 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch D

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 30 l3-interface 30
admin@XorPlus# set vlans vlan-id 50 l3-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
```

```

Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them IP addresses.

```

admin@XorPlus# set vlan-interface interface 30 vif 40 address 192.168.30.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure an IBGP connection.

```

admin@XorPlus# set protocols bgp bgp-id 4.4.4.4
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.30.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.30.1 local-ip 192.168.30.2
admin@XorPlus# set protocols bgp peer 192.168.30.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.50.2 local-ip 192.168.50.1
admin@XorPlus# set protocols bgp peer 192.168.50.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Viewing BGP Peer Status on Switch B

```

admin@XorPlus# run show bgp peers detail
Peer 2: local 192.168.10.2/179 remote 192.168.10.1/179
Peer ID: 1.1.1.1
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
Peer 2: local 192.168.30.1/179 remote 192.168.30.2/179
Peer ID: 4.4.4.4
Peer State: ESTABLISHED

```

```

Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
Peer 3: local 192.168.40.1/179 remote 192.168.40.2/179
Peer ID: 3.3.3.3
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#

```

Configuring Switch A to Advertise Route 192.168.10.0/24

```

admin@XorPlus# set policy policy-statement direct-to-bgp term t1 from protocol connected
admin@XorPlus# set policy policy-statement direct-to-bgp term t1 from network4
192.168.10.0/24
admin@XorPlus# set policy policy-statement direct-to-bgp term t1 then accept
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols bgp export direct-to-bgp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

View the BGP routing table of Switch B:

```

admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
*> 192.168.10.0/24 192.168.20.1 1.1.1.1 100 ?
admin@XorPlus#

```

View the BGP routing table of Switch C:

```

admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
192.168.10.0/24 192.168.20.1 2.2.2.2 100 ?
admin@XorPlus#

```

The preceding command output display that the route to destination 192.168.10.0/24 becomes invalid because the next hop address of this route is unreachable.

Configuring Switch B to Advertise a Connected Route

```

admin@XorPlus# set policy policy-statement direct-to-bgp term t1 from protocol connected
admin@XorPlus# set policy policy-statement direct-to-bgp term t1 then accept
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols bgp export direct-to-bgp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Ping 192.168.10.1 on Switch C:

```

admin@XorPlus# run ping 192.168.10.1
PING 192.168.10.1 (192.168.10.1) 56(84) bytes of data.
64 bytes from 192.168.10.1: icmp_req=1 ttl=63 time=4.68 ms
64 bytes from 192.168.10.1: icmp_req=2 ttl=63 time=4.46 ms
64 bytes from 192.168.10.1: icmp_req=3 ttl=63 time=5.35 ms
64 bytes from 192.168.10.1: icmp_req=4 ttl=63 time=4.52 ms
64 bytes from 192.168.10.1: icmp_req=5 ttl=63 time=4.51 ms
192.168.10.1 ping statistics -
5 packets transmitted, 5 received, 0% packet loss, time 4017ms
rtt min/avg/max/mdev = 4.460/4.709/5.358/0.338 ms
admin@XorPlus#

```

View the BGP routing table of Switch C:

```

admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
192.168.10.0/24 192.168.20.1 2.2.2.2 100 ?
*> 192.168.20.0/24 192.168.40.1 2.2.2.2?
*>192.168.30.0/24 192.168.40.1 2.2.2.2?
admin@XorPlus#

```

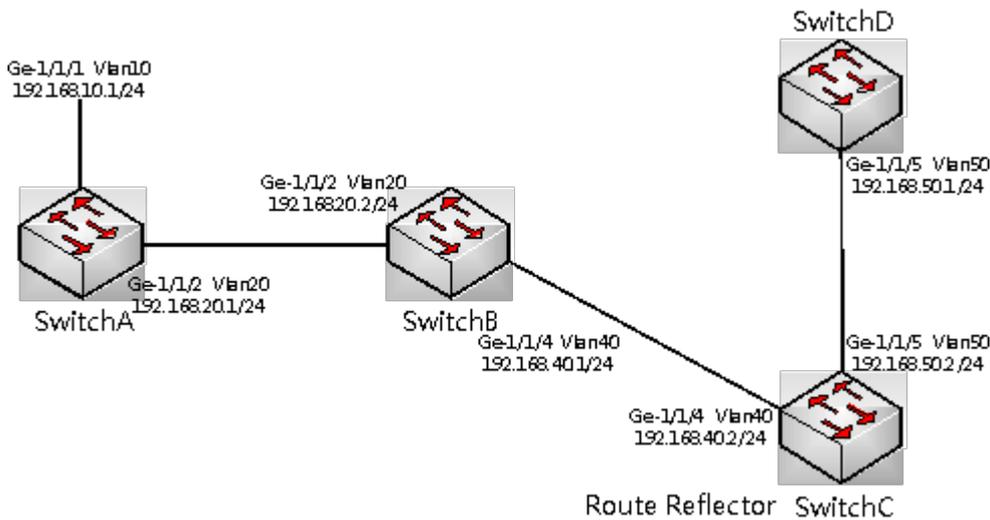
BGP Route Reflector Configuration Example

The IBGP network should be formed without interrupting fullmesh BGP connections between Switch B, Switch C, and Switch D, and call for simplified device configuration and management.

Configure Switch B, Switch C, and Switch D to have IBGP connections. Between Switch A and Switch B should be an EBGP connection.

Configure Switch C as a route reflector with clients Switch B and Switch D.

Figure 5-10. BGP route reflector



Configuring Switch A

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 10 13-interface 10
admin@XorPlus# set vlans vlan-id 20 13-interface 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them each IP addresses.

```
admin@XorPlus# set vlan-interface interface 10 vif 10 address 192.168.10.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP connection.

```

admin@XorPlus# set protocols bgp bgp-id 1.1.1.1
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.20.2 local-ip 192.168.20.1
admin@XorPlus# set protocols bgp peer 192.168.20.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B

Configure the VLAN that each interface belongs to.

```

admin@XorPlus# set vlans vlan-id 20 l3-interface 20
admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them each IP addresses.

```

admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure EBGP and IBGP connections.

```

admin@XorPlus# set policy policy-statement p2 term t1 from protocol bgp
admin@XorPlus#
admin@XorPlus# set policy policy-statement p2 term t1 from network4 192.168.10.0/24
admin@XorPlus# set policy policy-statement p2 term t1 then nexthop4 192.168.40.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols bgp bgp-id 2.2.2.2
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 100
admin@XorPlus# set protocols bgp peer 192.168.20.1 local-ip 192.168.20.2
admin@XorPlus# set protocols bgp peer 192.168.20.1 next-hop-self true

```

```

admin@XorPlus# set protocols bgp peer 192.168.40.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.40.2 local-ip 192.168.40.1
admin@XorPlus# set protocols bgp peer 192.168.40.2 next-hop-self true
admin@XorPlus# set protocols bgp export p2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure the VLAN that each interface belongs to.

```

admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# set vlans vlan-id 50 l3-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them each IP addresses.

```

admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.2 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure an IBGP connection.

```

admin@XorPlus# set protocols bgp bgp-id 3.3.3.3
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.40.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.40.1 local-ip 192.168.40.2
admin@XorPlus# set protocols bgp peer 192.168.40.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.50.1 local-ip 192.168.50.2
admin@XorPlus# set protocols bgp peer 192.168.50.1 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure IBGP connections for the route reflector clients.

```
admin@XorPlus# set protocols bgp route-reflector cluster-id 3.3.3.3
admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp peer 192.168.40.1 client true
admin@XorPlus# set protocols bgp peer 192.168.50.1 client true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch D

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 30 l3-interface 30
admin@XorPlus# set vlans vlan-id 50 l3-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them each IP addresses.

```
admin@XorPlus# set vlan-interface interface 30 vif 30 address 192.168.30.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an IBGP connection.

```
admin@XorPlus# set protocols bgp bgp-id 4.4.4.4
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.30.1 as 200
admin@XorPlus# set protocols bgp peer 192.168.30.1 local-ip 192.168.30.2
admin@XorPlus# set protocols bgp peer 192.168.30.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.50.2 local-ip 192.168.50.1
admin@XorPlus# set protocols bgp peer 192.168.50.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Verifying Configurations

View the BGP routing table of Switch B:

```
admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
*> 192.168.10.0/24 192.168.20.11.1.1.1200 ?
admin@XorPlus#
```

View the BGP routing table of Switch D:

```
admin@XorPlus# run show bgp routes detail
192.168.10.0/24
From peer: 3.3.3.3
Route: Not Used
Origin: INCOMPLETE
AS Path: 200
Nexthop: 192.168.40.1
Multiple Exit Discriminator: 0
Local Preference: 100
Originator ID: 2.2.2.2
Cluster List: 3.3.3.3
```

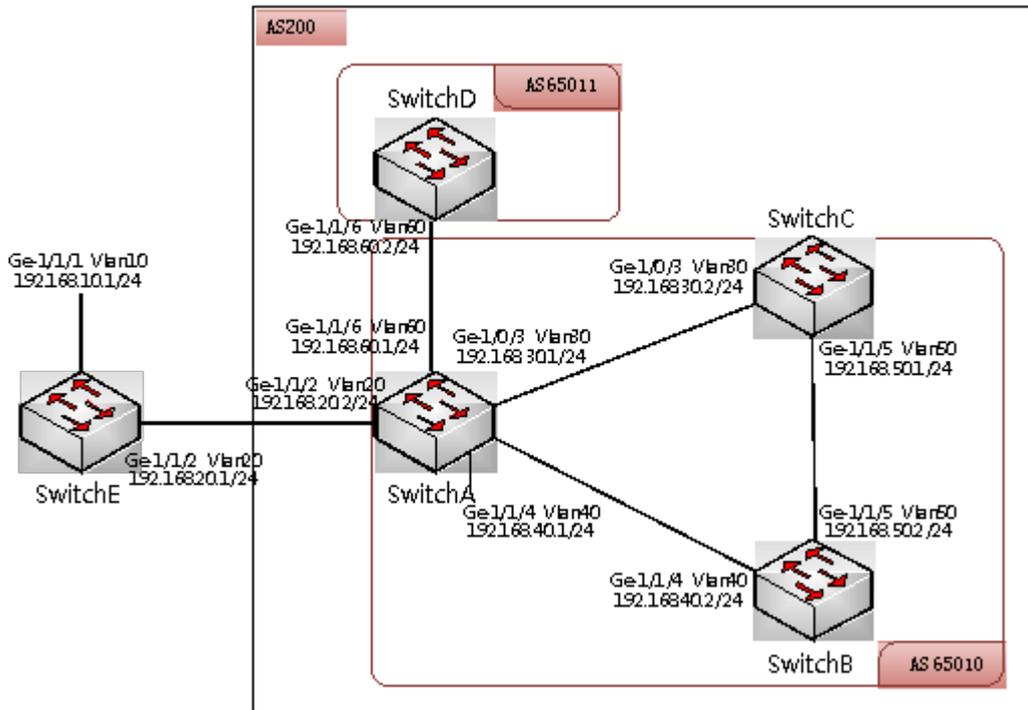
BGP Confederation Configuration Example

Configure a BGP confederation on each switch in AS 200 to divide AS 200 into two sub-AS's: AS 65010 and AS 65011. To reduce the number of IBGP connections, three switches in AS 65010 establish fullmesh IBGP connections.

Configure BGP confederation members Switch A, Switch B, Switch C, and Switch D. Between Switch A and Switch D is an EBGP connection within AS 200.

Configure Switch A to connect without AS 200 to Switch E.

Figure 5-11. BGP confederation configuration



Configuring Switch A

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 20 l3-interface 20
admin@XorPlus# set vlans vlan-id 30 l3-interface 30
admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# set vlans vlan-id 60 l3-interface 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# set interface gigabit-ethernet ge-1/1/6 family ethernet-switching
native-vlan-id 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them each IP addresses.

```
admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 30 vif 30 address 192.168.30.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 60 vif 60 address 192.168.60.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
```

```
Commit OK.
Save done.
admin@XorPlus#
```

Configure EBGP and IBGP connections within confederation AS 200.

```
admin@XorPlus# set protocols bgp bgp-id 2.2.2.2
admin@XorPlus# set protocols bgp local-as 65010
admin@XorPlus# protocols bgp confederation identifier 200
admin@XorPlus# protocols bgp confederation disable false
admin@XorPlus# set protocols bgp peer 192.168.30.2 as 65010
admin@XorPlus# set protocols bgp peer 192.168.30.2 local-ip 192.168.30.1
admin@XorPlus# set protocols bgp peer 192.168.30.2 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.30.2 confederation-member true
admin@XorPlus# set protocols bgp peer 192.168.40.2 as 65010
admin@XorPlus# set protocols bgp peer 192.168.40.2 local-ip 192.168.40.1
admin@XorPlus# set protocols bgp peer 192.168.40.2 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.40.2 confederation-member true
admin@XorPlus# set protocols bgp peer 192.168.60.2 as 65011
admin@XorPlus# set protocols bgp peer 192.168.60.2 local-ip 192.168.60.1
admin@XorPlus# set protocols bgp peer 192.168.60.2 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.60.2 confederation-member true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP connection without confederation AS 200.

```
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 100
admin@XorPlus# set protocols bgp peer 192.168.20.1 local-ip 192.168.20.2
admin@XorPlus# set protocols bgp peer 192.168.20.1 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch B

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 40 13-interface 40
admin@XorPlus# set vlans vlan-id 50 13-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/4 family ethernet-switching
native-vlan-id 40
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them each IP addresses.

```

admin@XorPlus# set vlan-interface interface 40 vif 40 address 192.168.40.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.2 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure an IBGP connection within confederation AS 200.

```

admin@XorPlus# set protocols bgp bgp-id 3.3.3.3
admin@XorPlus# set protocols bgp local-as 65010
admin@XorPlus# protocols bgp confederation identifier 200
admin@XorPlus# protocols bgp confederation disable false
admin@XorPlus# set protocols bgp peer 192.168.40.1 as 65010
admin@XorPlus# set protocols bgp peer 192.168.40.1 local-ip 192.168.40.2
admin@XorPlus# set protocols bgp peer 192.168.40.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.40.1 confederation-member true
admin@XorPlus# set protocols bgp peer 192.168.50.1 as 65010
admin@XorPlus# set protocols bgp peer 192.168.50.1 local-ip 192.168.50.2
admin@XorPlus# set protocols bgp peer 192.168.50.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.1 confederation-member true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C**Configure the VLAN that each interface belongs to.**

```

admin@XorPlus# set vlans vlan-id 30 13-interface 30
admin@XorPlus# set vlans vlan-id 50 13-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/5 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces and assign them each IP addresses.

```

admin@XorPlus# set vlan-interface interface 30 vif 40 address 192.168.30.2 prefix-length 24
admin@XorPlus# set vlan-interface interface 50 vif 50 address 192.168.50.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.

```

```
Commit OK.
Save done.
admin@XorPlus#
```

Configure an IBGP connection within confederation AS 200.

```
admin@XorPlus# set protocols bgp bgp-id 4.4.4.4
admin@XorPlus# set protocols bgp local-as 65010
admin@XorPlus# protocols bgp confederation identifier 200
admin@XorPlus# protocols bgp confederation disable false
admin@XorPlus# set protocols bgp peer 192.168.30.1 as 65010
admin@XorPlus# set protocols bgp peer 192.168.30.1 local-ip 192.168.30.2
admin@XorPlus# set protocols bgp peer 192.168.30.1 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.30.1 confederation-member true
admin@XorPlus# set protocols bgp peer 192.168.50.2 as 65010
admin@XorPlus# set protocols bgp peer 192.168.50.2 local-ip 192.168.50.1
admin@XorPlus# set protocols bgp peer 192.168.50.2 next-hop-self true
admin@XorPlus# set protocols bgp peer 192.168.50.2 confederation-member true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch D

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 60 l3-interface 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/6 family ethernet-switching
native-vlan-id 60
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them each IP addresses.

```
admin@XorPlus# set vlan-interface interface 60 vif 60 address 192.168.60.2 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP connection within confederation AS 200.

```
admin@XorPlus# set protocols bgp bgp-id 5.5.5.5
admin@XorPlus# set protocols bgp local-as 65011
admin@XorPlus# protocols bgp confederation identifier 200
admin@XorPlus# protocols bgp confederation disable false
admin@XorPlus# set protocols bgp peer 192.168.60.2 as 65010
admin@XorPlus# set protocols bgp peer 192.168.60.2 local-ip 192.168.60.1
admin@XorPlus# set protocols bgp peer 192.168.60.2 next-hop-self true
```

```
admin@XorPlus# set protocols bgp peer 192.168.60.2 confederation-member true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch E

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 10 l3-interface 10
admin@XorPlus# set vlans vlan-id 20 l3-interface 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them IP addresses.

```
admin@XorPlus# set vlan-interface interface 10 vif 10 address 192.168.10.1 prefix-length 24
admin@XorPlus# set vlan-interface interface 20 vif 20 address 192.168.20.1 prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP connection.

```
admin@XorPlus# set protocols bgp bgp-id 1.1.1.1
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.20.2 local-ip 192.168.20.1
admin@XorPlus# set protocols bgp peer 192.168.20.2 next-hop-self true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Verifying the Configuration

View the BGP routing table of Switch B:

```

admin@XorPlus# run show bgp routes detail
192.168.10.0/24
From peer: 2.2.2.2
Route: Not Used
Origin: INCOMPLETE
AS Path: 100
Nextthop: 192.168.20.1
Multiple Exit Discriminator: 0
Local Preference: 100
View the BGP routing table of Switch D:
admin@XorPlus# run show bgp routes detail
192.168.10.0/24
From peer: 15.15.15.15
Route: Not Used
Origin: INCOMPLETE
AS Path: (65010) 100
Nextthop: 192.168.30.2
Local Preference: 100

```

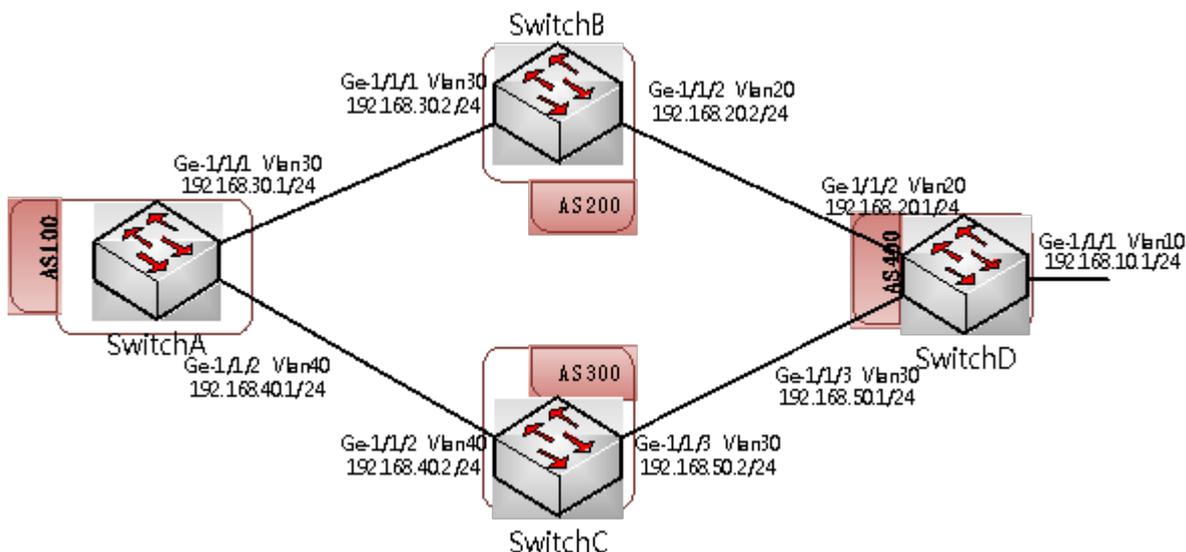
BGP Load Balancing Configuration Example

Configure load balancing on Switch A.

Configure EBGP connections between Switch B and Switch A and between Switch B and Switch D.

Configure EBGP connections between Switch C and Switch A and between Switch C and Switch D.

Figure 5-12. BGP load balancing



Configuring Switch A

Configure the VLAN that each interface belongs to.

```

admin@XorPlus# set vlans vlan-id 30 13-interface vlan30
admin@XorPlus# set vlans vlan-id 40 13-interface vlan40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching

```

```

native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 40
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them each IP addresses.

```

admin@XorPlus# set vlan-interface interface vlan30 vif vlan30 address 192.168.30.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan40 vif vlan40 address 192.168.40.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure an EBGp connection.

```

admin@XorPlus# set protocols bgp bgp-id 1.1.1.1
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 192.168.30.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.30.2 local-ip 192.168.30.1
admin@XorPlus# set protocols bgp peer 192.168.40.2 as 300
admin@XorPlus# set protocols bgp peer 192.168.40.2 local-ip 192.168.30.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch B

Configure the VLAN that each interface belongs to.

```

admin@XorPlus# set vlans vlan-id 20 l3-interface vlan20
admin@XorPlus# set vlans vlan-id 30 l3-interface vlan30
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 30
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them IP addresses.

```

admin@XorPlus# set vlan-interface interface vlan20 vif vlan20 address 192.168.20.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan30 vif vlan30 address 192.168.30.2
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure an EBGP connection.

```

admin@XorPlus# set protocols bgp bgp-id 2.2.2.2
admin@XorPlus# set protocols bgp local-as 200
admin@XorPlus# set protocols bgp peer 192.168.20.1 as 400
admin@XorPlus# set protocols bgp peer 192.168.20.1 local-ip 192.168.20.2
admin@XorPlus# set protocols bgp peer 192.168.30.1 as 100
admin@XorPlus# set protocols bgp peer 192.168.30.2 local-ip 192.168.30.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configuring Switch C

Configure the VLAN that each interface belongs to.

```

admin@XorPlus# set vlans vlan-id 40 l3-interface 40
admin@XorPlus# set vlans vlan-id 50 l3-interface 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 40
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure the VLAN interfaces, and assign them IP addresses.

```

admin@XorPlus# set vlan-interface interface vlan40 vif vlan40 address 192.168.40.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan50 vif vlan50 address 192.168.50.2
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
Configure an EBGP connection.
admin@XorPlus# set protocols bgp bgp-id 3.3.3.3
admin@XorPlus# set protocols bgp local-as 300
admin@XorPlus# set protocols bgp peer 192.168.40.1 as 100
admin@XorPlus# set protocols bgp peer 192.168.40.1 local-ip 192.168.40.2

```

```
admin@XorPlus# set protocols bgp peer 192.168.50.1 as 400
admin@XorPlus# set protocols bgp peer 192.168.50.1 local-ip 192.168.50.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring Switch D

Configure the VLAN that each interface belongs to.

```
admin@XorPlus# set vlans vlan-id 10 l3-interface vlan10
admin@XorPlus# set vlans vlan-id 20 l3-interface vlan20
admin@XorPlus# set vlans vlan-id 50 l3-interface vlan50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 10
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 20
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 family ethernet-switching
native-vlan-id 50
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure the VLAN interfaces, and assign them IP addresses.

```
admin@XorPlus# set vlan-interface interface vlan10 vif vlan10 address 192.168.10.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan20 vif vlan20 address 192.168.20.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan50 vif vlan50 address 192.168.50.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configure an EBGP Connection

```
admin@XorPlus# set protocols bgp bgp-id 4.4.4.4
admin@XorPlus# set protocols bgp local-as 400
admin@XorPlus# set protocols bgp peer 192.168.20.2 as 200
admin@XorPlus# set protocols bgp peer 192.168.20.2 local-ip 192.168.20.1
admin@XorPlus# set protocols bgp peer 192.168.50.2 as 300
admin@XorPlus# set protocols bgp peer 192.168.50.2 local-ip 192.168.50.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Viewing BGP Peer Status on Switch B

```
admin@XorPlus# run show bgp peers detail
Peer 2: local 192.168.20.2/179 remote 192.168.20.1/39912
Peer ID: 1.1.1.1
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
Peer 2: local 192.168.30.2/16808 remote 192.168.30.1/179
Peer ID: 4.4.4.4
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Viewing BGP Peer Status on Switch C

```
admin@XorPlus# run show bgp peers detail
Peer 2: local 192.168.40.2/179 remote 192.168.40.1/38815
Peer ID: 1.1.1.1
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
Peer 2: local 192.168.50.2/49923 remote 192.168.50.1/179
Peer ID: 4.4.4.4
Peer State: ESTABLISHED
Admin State: START
```

```

Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 20, Updates Sent: 2
Messages Received: 634, Messages Sent: 611
Time since last received update: 1685 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 15995 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Configuring Switch D to Advertise Route 192.168.10.0/24

Configure Switch A to enable BGP multipath:

```

admin@XorPlus# set protocols bgp multipath disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

Configure Switch D to advertise route 192.168.10.0/24:

```

admin@XorPlus# set policy policy-statement direct-to-bgp term t1 from protocol connected
admin@XorPlus# set policy policy-statement direct-to-bgp term t1 from network4
192.168.10.0/24
admin@XorPlus# set policy policy-statement direct-to-bgp term t1 then accept
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set protocols bgp export direct-to-bgp
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

View the BGP routing table of Switch A:

```

admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
*> 192.168.10.0/24 192.168.30.2 2.2.2.2 200 400 ?
192.168.10.0/24 192.168.40.2 3.3.3.3 200 400 ?
admin@XorPlus#

```

As expected, Switch A is not load balancing because it does not view the paths as "equal," but as different AS paths.

Configuring BGP Multipath Path-Relax on Switch A

```
admin@XorPlus# set protocols bgp multipath path-relax true
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

View the BGP routing table and IP routing table of Switch A:

```
admin@XorPlus# run show bgp routes
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix Nexthop Peer AS Path
-----
*> 192.168.10.0/24 192.168.30.2 2.2.2.2 200 400 ?
*> 192.168.10.0/24 192.168.40.2 3.3.3.3 200 400 ?
admin@XorPlus#
admin@XorPlus# run show route table ipv4 unicast ebgp
192.168.10.0/24 [ebgp(20)/0]
> to 192.168.30.2 via vlan30/vlan30
192.168.10.0/24 [ebgp(20)/0]
> to 192.168.40.2 via vlan40/vlan40
admin@XorPlus#
admin@XorPlus#
```

IPv6 BGP Configuration

IPv6 BGP configuration

- IPv6 BGP Introduction
- Building Peering Sessions
 - EBGp Peering
 - IBGP Peering
- Establish BGP Peer Use 4-byte-AS-Number
- Sources of Routing Updates
 - Injecting Information Dynamically into BGP
 - Injecting Information Statically into BGP
- BGP Attributes
 - The NEXT_HOP Attribute
 - The AS_PATH Attribute
 - The LOCAL_PREF Attribute
 - The MULTI_EXIT_DISC Attribute
 - The COMMUNITY Attribute
- BGP-4 Aggregation
- Synchronization

- Controlling Large-Scale Autonomous System
 - Confederations
 - Route Reflectors
- Redundancy and Load Balancing
- Designing Stable Internets

IPv6 BGP Introduction

BGP is a path vector protocol used to carry routing information between autonomous systems. The term path vector comes from the fact that BGP routing information carries a sequence of AS numbers that identifies the path of AS's that a network prefix has traversed. The path information associated with the prefix is used to enable loop prevention.

BGP uses TCP as its transport protocol (port 179). This ensures that all the transport reliability (such as retransmission) is taken care of by TCP and does not need to be implemented in BGP, thereby simplifying the complexity associated with designing reliability into the protocol itself.

Routers that run a BGP routing process are often referred to as BGP speakers. Two BGP speakers that form a TCP connection between one another for the purpose of exchanging routing information, are referred to as neighbors or peers. Peer routers exchange open messages to determine the connection parameters. These messages are used to communicate values such as the BGP speaker's version number.

BGP also provides a mechanism to gracefully close a connection with a peer. In other words, in the event of a disagreement between the peers, be it resultant of configuration, incompatibility, operator intervention, or other circumstances, a NOTIFICATION error message is sent, and the peer connection does not get established or is torn down if it's already established. The benefit of this mechanism is that both peers understand that the connection could not be established or maintained and do not waste resources that would otherwise be required to maintain or blindly reattempt to establish the connection. The graceful close mechanism simply ensures that all outstanding messages, primarily NOTIFICATION error messages, are delivered before the TCP session is closed.

Initially, when a BGP session is established between a set of BGP speakers, all candidate BGP routes are exchanged. After the session has been established and the initial route exchange has occurred, only incremental updates are sent as network information changes.

Routes are advertised between a pair of BGP routers in UPDATE messages. The UPDATE message contains, among other things, a list of <length, prefix> tuples that indicate the list of destinations that can be reached via a BGP speaker. The UPDATE message also contains the path attributes, which include such information as the degree of preference for a particular route and the list of AS's that the route has traversed.

In the event that a route becomes unreachable, a BGP speaker informs its neighbors by withdrawing the invalid route. Withdrawn routes are part of the UPDATE message. These routes are no longer available for use. If information associated with a route has changed or a new path for the same prefix has been selected, a withdrawal is not required; it is enough to just advertise a replacement route.

If no routing changes occur, the routers exchange only KEEPALIVE packets.

KEEPALIVE messages are sent periodically between BGP neighbors to ensure that the connection is kept alive. KEEPALIVE packets (19 bytes each) should not cause any strain on the router CPU or link bandwidth because they consume a minimal amount of bandwidth.

Building Peering Sessions

Command References

```
admin@XorPlus# set protocols bgp bgp-id <ipv4 address>
```

Note: This command is to configure a BGP-ID. It indicates a BGP router uniquely.

```
admin@XorPlus# set protocols bgp local-as <as-number>
```

Note: This command is to configure an AS-number. It tells us which AS the BGP router is in. It should be expressed using a two-byte length digit or a 4-byte length digit <0-4294967295> or <0..65535>.<0..65535>.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> as <peer-as-number>
```

Note: This command is to specify a BGP peer and it's AS-number.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> local-ip <local-ipv6-address>
```

Note: This command is to specify the local IPv6 address for BGP peer session.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> ipv6-unicast <true/false>
```

Note: IPv6 BGP route will be propagated to its BGP peer after enabling IPv6-unicast.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> ipv4-unicast <true/false>
```

Note: The IPv4 BGP route entry should be propagated via BGP update packet. IPv4-unicast was disabled by default.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> default-route-advertise disable  
<true/false>
```

Note: This command is to advertise a default route entry to its BGP peer after user enables default-route-advertise.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> allow-as-loop <true/false>
```

Note: This command is to receive the BGP route entry with its own as number. By default, the BGP route entry with its own AS number will be dropped to prevent loops.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> holdtime <0 or 3-65535>
```

Note: This command is to set the hold-timer value, and the keepalive timer will be holdtime/3, holdtime should be 0. This means, the BGP peering session will not expire anymore, and holdtime should be the digit 3-65535.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> md5-password <text>
```

Note: This command is to configure a md5-password. the two peers must have the same md5-password.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> next-hop-self <true/false>
```

Note: This command is to enable next-hop-self when user configures IBGP peer session. The next-hop will be set its own IPv6 address after user enables next-hop-self.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> prefix-limit maximum <1-12000>
```

Note: This command is to set the maximum BGP route entry limit for one BGP peer.

```
admin@XorPlus# set protocols bgp peer <peer-ipv6-address> public-as-only <true/false>
```

Note: The private AS-number will be removed from the AS-path if the private AS number in the AS-path after user enables public-as-only.

- EBGPeering
- IBGP Peering

EBGP Peering

EBGP Peering:

Directly connected peer

Non-Directly connected peer

Establish ebgp peer use loopback interface

(1) Directly connected peer

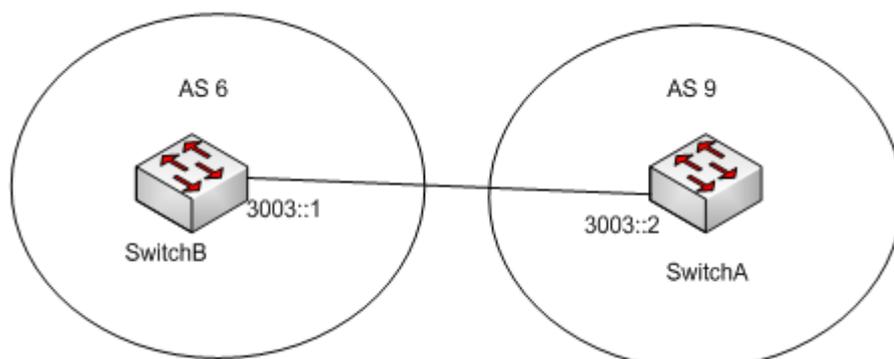


Figure 1-4

Step 1: Configure bgp-id and local-as

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
```

Step 2: Configure bgp peer

Switch A:

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "6"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 3: Enable ipv6-unicast

Switch A:

```
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 4: Check bgp peer status:

Switch A:

```
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/53149
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 3, Messages Sent: 3
Time since last received update: n/a
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 33 seconds
```

```

Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/53149 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 0
Messages Received: 3, Messages Sent: 3
Time since last received update: n/a
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 41 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

(2) Establish EBGP via non-direct-connected interfaces

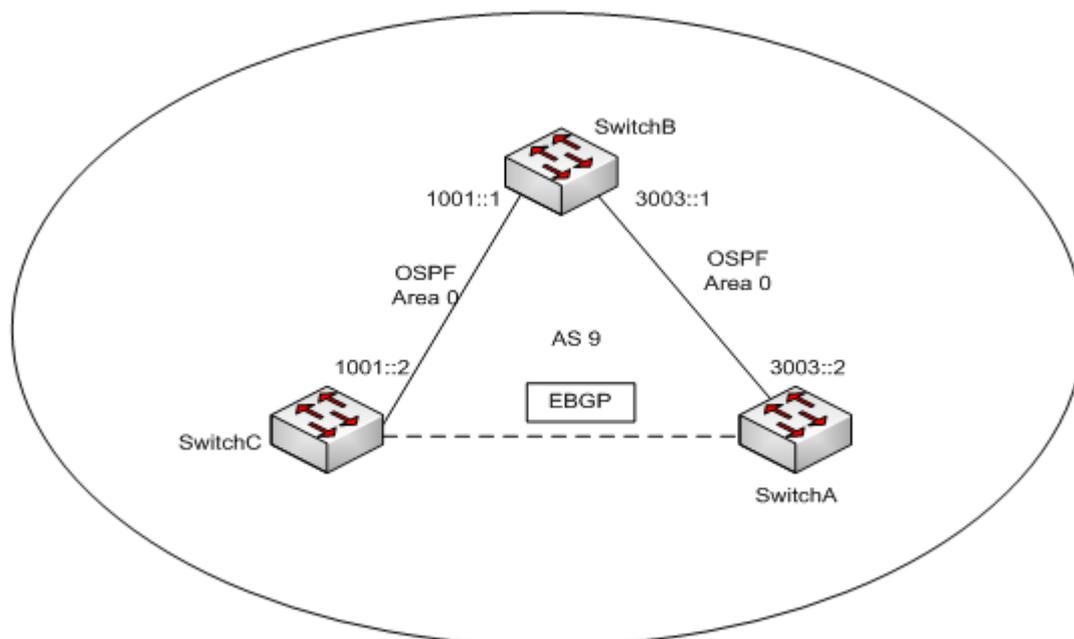


Figure 1-5

Step 1: SwitchA SwitchB SwitchC Enable OSPFV3

Switch A:

```

admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 9.9.9.9
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::1
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan300 vif vlan300 address
3003::2

```

Switch B:

```
admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 6.6.6.6
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan100 vif vlan100 address
1001::1
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan300 vif vlan300 address
3003::1
```

Switch C:

```
admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 26.26.26.26
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan100 vif vlan100 address
1001::2
```

Step 2: Check ospfv3 status on SwitchB

```
admin@XorPlus# run show ospf6 neighbor
Address                               Interface                               State   Router ID
Pri    Dead
-----
fe80::ca0a:a9ff:204:4928             vlan100/vlan100                       Full    26.26.26.26
128    35
fe80::ca0a:a9ff:5ae:a66             vlan300/vlan300                       Full    9.9.9.9
128    38
```

Note: Two ospf6 neighbor all established

Step 3: Check the route table on SwitchA SwitchC

Note: there should be 3003::/64 route entry on SwitchC, and 1001::/64 should in SwitchA's route table:

Switch A:

```
admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                                Port
-----
2001::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
4001::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
3003::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
1001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF te-1/1/46
6006::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
5005::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
1001::                                     ffff:ff00::
04:7D:7B:62:93:FF te-1/1/46
```

Switch C:

```

admin@XorPlus# run show route forward-route ipv6 all
Destination                                     NetMask
NextHopMac           Port
-----
3003::               ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   te-1/1/49
6006::               ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28   connected
2002::               ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28   connected
5005::               ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   te-1/1/49
1001::               ffff:ff00::
C8:0A:A9:04:49:28   connected

```

Step 4: Configure EBGP on SwitchA SwitchC

Switch A:

```

admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 1001::2 local-ip "3003::2"
admin@XorPlus#set protocols bgp peer 1001::2 as "26"
admin@XorPlus#set protocols bgp peer 1001::2 ipv6-unicast true

```

Switch C:

```

admin@XorPlus#set protocols bgp bgp-id 26.26.26.26
admin@XorPlus#set protocols bgp local-as "26"
admin@XorPlus#set protocols bgp peer 3003::2 local-ip "1001::2"
admin@XorPlus#set protocols bgp peer 3003::2 as "9"
admin@XorPlus#set protocols bgp peer 3003::2 ipv6-unicast true

```

Step 5: Check bgp peer status

```

admin@XorPlus# run show bgp peers detail
Peer 1: local 3003::2/60737 remote 1001::2/179
Peer ID: 26.26.26.26
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 26
Updates Received: 0, Updates Sent: 0
Messages Received: 2, Messages Sent: 2
Time since last received update: n/a
Number of transitions to ESTABLISHED: 2
Time since last entering ESTABLISHED state: 13 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

(3) Establish ebgp peer use loopback interface

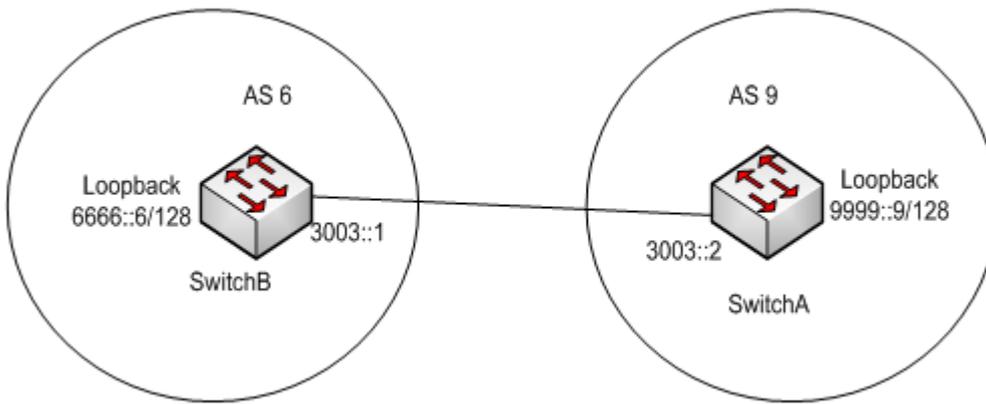


Figure 1-6

Step 1: Configure loopback interface

Switch A:

```
admin@XorPlus# set vlan-interface loopback address 6666::6 prefix-length 128
```

Switch B:

```
admin@XorPlus# set vlan-interface loopback address 9999::9 prefix-length 128
```

Step 2: Configure static route on SwitchA SwitchB

Switch A:

```
admin@XorPlus#set protocols static route 9999::9/128 next-hop 3003::2
admin@XorPlus# run show route forward-host ipv6 all
```

Address	HWaddress	Port
9999::9	C8:0A:A9:AE:0A:66	te-1/1/46
3003::2	C8:0A:A9:AE:0A:66	te-1/1/46
6666::6	04:7D:7B:62:93:FF	connected

Switch B:

```
admin@XorPlus# set protocols static route 6666::6/128 next-hop 3003::1
admin@XorPlus# run show route forward-host ipv6 all
```

Address	HWaddress	Port
3003::1	04:7D:7B:62:93:FF	te-1/1/46
9999::9	C8:0A:A9:AE:0A:66	connected
6666::6	04:7D:7B:62:93:FF	te-1/1/46

Step 3: Configure bgp-id and local-as

Note: The two parameters must commit at the same time it will commit failed if lack either of the two

Switch A:

```
admin@XorPlus#set protocols bgp bgp-id 6.6.6.6
admin@XorPlus#set protocols bgp local-as "6"
```

Switch B:

```
admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
```

Step 4: Configure one bgp peer

Switch A:

```
admin@XorPlus#set protocols bgp peer 9999::9 local-ip "6666::6"
admin@XorPlus#set protocols bgp peer 9999::9 as "9"
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 6666::6 local-ip "9999::9"
admin@XorPlus#set protocols bgp peer 6666::6 as "6"
```

Step 5: Enable ipv6-unicast

Switch A:

```
admin@XorPlus# set protocols bgp peer 9999::9 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 6666::6 ipv6-unicast true
```

Step 6: Check bgp peer status

Switch A:

```
admin@XorPlus# run show bgp peers detail 9999::9
Peer 1: local 6666::6/33573 remote 9999::9/179
  Peer ID: 9.9.9.9
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 9
  Updates Received: 1, Updates Sent: 0
  Messages Received: 48, Messages Sent: 48
  Time since last received update: 1218 seconds
  Number of transitions to ESTABLISHED: 4
  Time since last entering ESTABLISHED state: 1238 seconds
  Retry Interval: 120 seconds
  Hold Time: 90 seconds, Keep Alive Time: 30 seconds
  Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
  Minimum AS Origination Interval: 0 seconds
  Minimum Route Advertisement Interval: 0 seconds
```

Switch B:

```

admin@XorPlus# run show bgp peers detail 6666::6
Peer 1: local 9999::9/179 remote 6666::6/33573
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 1
Messages Received: 48, Messages Sent: 51
Time since last received update: n/a
Number of transitions to ESTABLISHED: 4
Time since last entering ESTABLISHED state: 1236 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

IBGP Peering

IBGP Peering:

Directly connected peer

Non-Directly connected peer

Establish bgp peer use loopback interface

(1) Directly connected peer

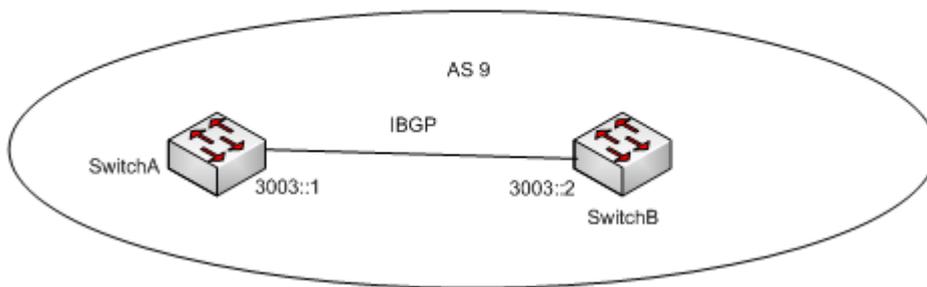


Figure 1-1

Step 1: Configure bgp-id and local-as

Note: The two parameters must commit at the same time it will commit failed if lack either of the two

Switch A:

```

admin@XorPlus#set protocols bgp bgp-id 6.6.6.6
admin@XorPlus#set protocols bgp local-as "9"

```

Switch B:

```

admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"

```

Step 2: Configure one bgp peer

Switch A:

```
admin@XorPlus#set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus#set protocols bgp peer 3003::2 as "9"
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus#set protocols bgp peer 3003::1 as "9"
```

Step 3: Enable ipv6-unicast

Switch A:

```
admin@XorPlus#set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 3003::1 ipv6-unicast true
```

Step 4: Enable next-hop-self

Switch A:

```
admin@XorPlus#set protocols bgp peer 3003::2 next-hop-self true
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 3003::1 next-hop-self true
```

Step 5: Check bgp peer status on SwitchA SwitchB

Switch A:

```
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/39351
  Peer ID: 9.9.9.9
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 9
  Updates Received: 1, Updates Sent: 0
  Messages Received: 27, Messages Sent: 26
  Time since last received update: 669 seconds
  Number of transitions to ESTABLISHED: 1
  Time since last entering ESTABLISHED state: 669 seconds
```

```

Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#

```

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/39351 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 1
Messages Received: 25, Messages Sent: 27
Time since last received update: n/a
Number of transitions to ESTABLISHED: 5
Time since last entering ESTABLISHED state: 643 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

(2) Non-Directly connected peer

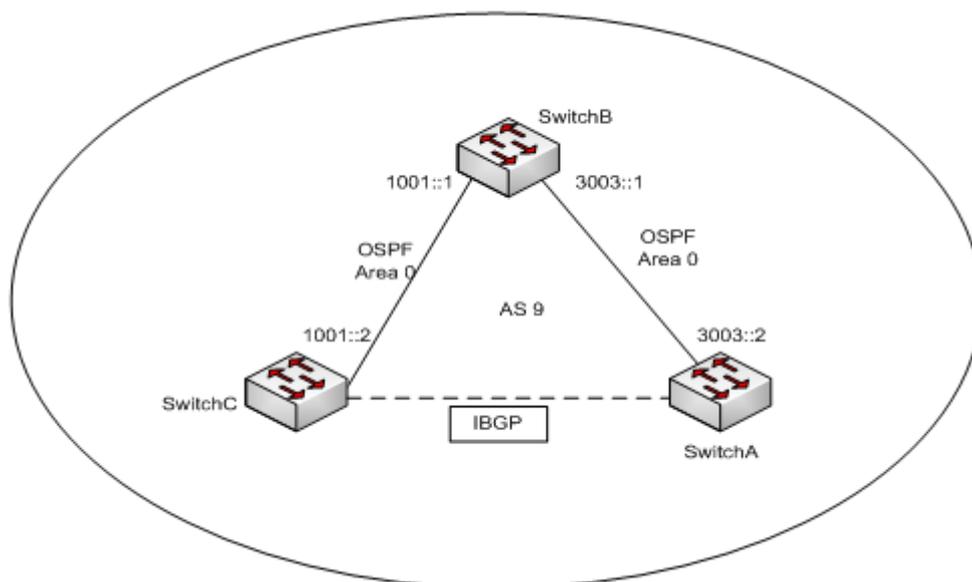


Figure 1-2

Step 1: SwitchA SwitchB SwitchC Enable OSPFV3

Switch A:

```

admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 9.9.9.9
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::1
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan300 vif vlan300 address
3003::2

```

Switch B:

```

admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 6.6.6.6
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vln100 vif vln100 address
1001::1
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vln300 vif vln300 address
3003::1

```

Switch C:

```

admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 26.26.26.26
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vln100 vif vln100 address
1001::2

```

Step 2: Check ospfv3 status on SwitchB

```

admin@XorPlus# run show ospf6 neighbor

```

Address	Interface	State	Router ID
Pri Dead			
-----	-----	-----	-----
-----	-----	-----	-----
fe80::ca0a:a9ff:204:4928 128 35	vln100/vln100	Full	26.26.26.26
fe80::ca0a:a9ff:5ae:a66 128 38	vln300/vln300	Full	9.9.9.9

Note: Two ospf6 neighbor all established

Step 3: Check the route table on SwitchA SwitchC

Note: there should be 3003::/64 route entry on SwitchC, and 1001::/64 should in SwitchA's route table:

Switch A:

```

admin@XorPlus# run show route forward-route ipv6 all

```

Destination	Port	NetMask
NextHopMac		
-----	-----	-----
-----	-----	-----
2001::		ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66	connected	
4001::		ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66	connected	
3003::		ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66	connected	
1001::		ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF	te-1/1/46	
6006::		ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66	connected	
5005::		ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66	connected	
1001::		ffff:ff00::
04:7D:7B:62:93:FF	te-1/1/46	

Switch C:

```

admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                                Port
-----
3003::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   te-1/1/49
6006::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28   connected
2002::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28   connected
5005::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   te-1/1/49
1001::                                     ffff:ff00::
C8:0A:A9:04:49:28   connected

```

Step 4: Configure IBGP on SwitchA SwitchC**Switch A:**

```

admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 1001::2 local-ip "3003::2"
admin@XorPlus#set protocols bgp peer 1001::2 as "9"
admin@XorPlus#set protocols bgp peer 1001::2 ipv6-unicast true

```

Switch C:

```

admin@XorPlus#set protocols bgp bgp-id 26.26.26.26
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 3003::2 local-ip "1001::2"
admin@XorPlus#set protocols bgp peer 3003::2 as "9"
admin@XorPlus#set protocols bgp peer 3003::2 ipv6-unicast true

```

Step 5: Check bgp peer status**Switch A:**

```

admin@XorPlus# run show bgp peers detail
Peer 1: local 3003::2/179 remote 1001::2/49225
  Peer ID: 26.26.26.26
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 9
  Updates Received: 0, Updates Sent: 0
  Messages Received: 2, Messages Sent: 2
  Time since last received update: n/a
  Number of transitions to ESTABLISHED: 1
  Time since last entering ESTABLISHED state: 23 seconds
  Retry Interval: 120 seconds
  Hold Time: 90 seconds, Keep Alive Time: 30 seconds
  Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
  Minimum AS Origination Interval: 0 seconds
  Minimum Route Advertisement Interval: 0 seconds

```

Note: As the peer 3003::2 could be connected via ospf6, there are route table for network 3003::/64

(3) Establish bgp peer use loopback interface

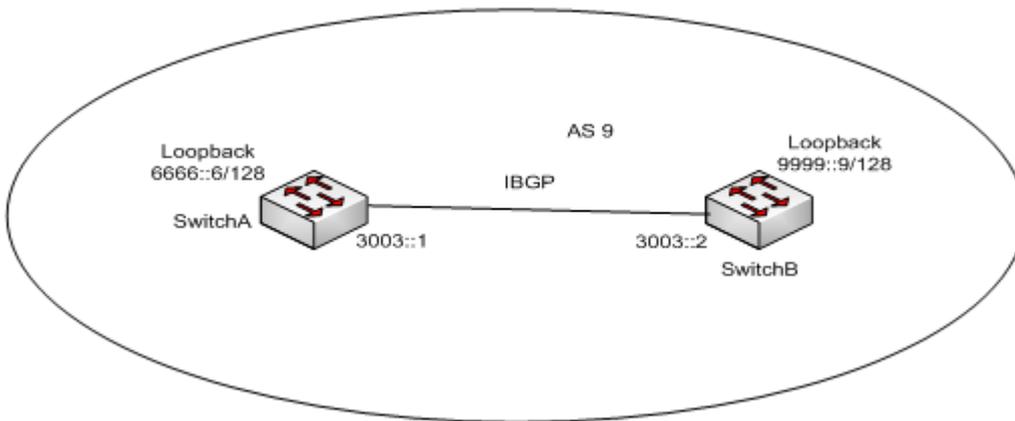


Figure 1-3

Step 1: Configure loopback interface

Switch A:

```
admin@XorPlus#set vlan-interface loopback address 6666::6 prefix-length 128
```

Switch B:

```
admin@XorPlus# set vlan-interface loopback address 9999::9 prefix-length 128
```

Step 2: Configure static route on SwitchA SwitchB

Switch A:

```
admin@XorPlus#set protocols static route 9999::9/128 next-hop 3003::2
admin@XorPlus# run show route forward-host ipv6 all
```

Address	HWaddress	Port
9999::9	C8:0A:A9:AE:0A:66	te-1/1/46
3003::2	C8:0A:A9:AE:0A:66	te-1/1/46
6666::6	04:7D:7B:62:93:FF	connected

Switch B:

```
admin@XorPlus# set protocols static route 6666::6/128 next-hop 3003::1
admin@XorPlus# run show route forward-host ipv6 all
```

Address	HWaddress	Port
3003::1	04:7D:7B:62:93:FF	te-1/1/46
9999::9	C8:0A:A9:AE:0A:66	connected
6666::6	04:7D:7B:62:93:FF	te-1/1/46

Step 3: Configure bgp-id and local-as

Note: The two parameters must commit at the same time it will commit failed if lack either of the two

Switch A:

```
admin@XorPlus#set protocols bgp bgp-id 6.6.6.6
admin@XorPlus#set protocols bgp local-as "9"
```

Switch B:

```
admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
```

Step 4: Configure bgp peer

Switch A:

```
admin@XorPlus#set protocols bgp peer 9999::9 local-ip "6666::6"
admin@XorPlus#set protocols bgp peer 9999::9 as "9"
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 6666::6 local-ip "9999::9"
admin@XorPlus#set protocols bgp peer 6666::6 as "9"
```

Step 5: Enable ipv6-unicast

Switch A:

```
admin@XorPlus# set protocols bgp peer 9999::9 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 6666::6 ipv6-unicast true
```

Step 6: Enable next-hop-self

Switch A:

```
admin@XorPlus#set protocols bgp peer 3003::2 next-hop-self true
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 3003::1 next-hop-self true
```

Step 7: check bgp peer status

Switch A:

```

admin@XorPlus# run show bgp peers detail 9999::9
Peer 1: local 6666::6/60097 remote 9999::9/179
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 3, Messages Sent: 3
Time since last received update: n/a
Number of transitions to ESTABLISHED: 3
Time since last entering ESTABLISHED state: 38 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#

```

Switch B:

```

admin@XorPlus# run show bgp peers detail 6666::6
Peer 1: local 9999::9/179 remote 6666::6/60097
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 4, Messages Sent: 5
Time since last received update: n/a
Number of transitions to ESTABLISHED: 3
Time since last entering ESTABLISHED state: 62 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

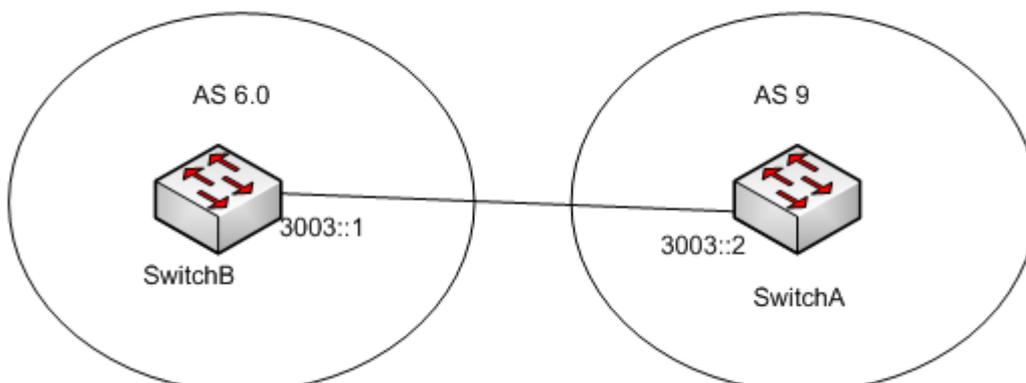
```

Establish BGP Peer Use 4-byte-AS-Number

Case 1

One Peer Enable 4-byte-as-number. Another Peer Didn't Enable 4-byte-as-number.

Figure 1-7



Step 1: Enable 4-Byte-AS-Number on Switch B

Switch B:

```
admin@XorPlus# set protocols bgp enable-4byte-as-numbers true
```

Step 2: Configure BGP-ID and Local-AS

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6.0"
```

Step 3: Configure BGP Peer

Note: User must use a special as number 23456 if local didn't enable 4-Byte-AS-Number but its peer already enabled 4-Byte-AS-Number.

Switch A:

```
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "23456"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 4: Check BFP Peer Status

Switch A:

```
admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/56968 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 23456
Updates Received: 0, Updates Sent: 0
Messages Received: 5, Messages Sent: 5
Time since last received update: n/a
Number of transitions to ESTABLISHED: 5
Time since last entering ESTABLISHED state: 87 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/56968
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 5, Messages Sent: 6
Time since last received update: n/a
Number of transitions to ESTABLISHED: 5
Time since last entering ESTABLISHED state: 100 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 5: Switch B Distribute a BGP route entry, then check the BGP route table on Switch A

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

Prefix                Nexthop                Peer                AS Path
-----
*> 9999::9/128        3003::2                0.0.0.0             ?
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
  From peer: 0.0.0.0
  Route: Winner
  Origin: INCOMPLETE
  AS Path:
  Nexthop: 3003::2
  Local Preference: 100
admin@XorPlus#

```

Switch A BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

Prefix                Nexthop                Peer                AS Path
-----
*> 9999::9/128        3003::1                6.6.6.6             23456 ?
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
  From peer: 6.6.6.6
  Route: Winner
  Origin: INCOMPLETE
  AS Path: 23456
  Nexthop: 3003::1
  Multiple Exit Discriminator: 0
  Local Preference: 100
  AS4 Path: 6.0
admin@XorPlus#

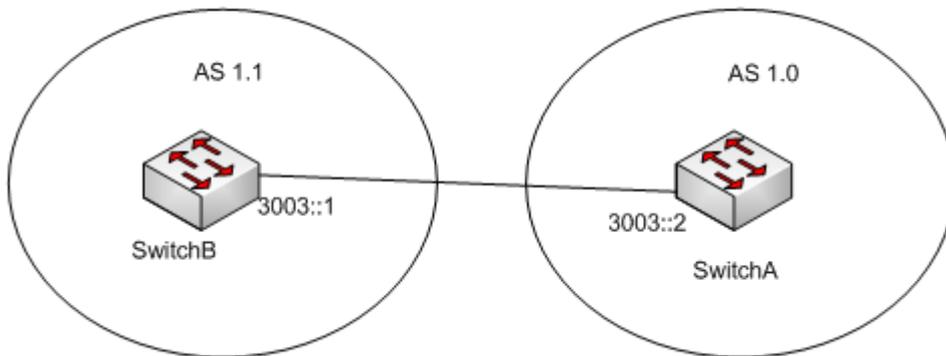
```

Note: User can see that the AS-path is 23456. The BGP route entry came from Switch B, as Switch A didn't support 4-Byte-AS-Number.

Case 2

Two Peer All Enable 4-byte-as-number

Figure 1-8



Step 1: Enable 4-Byte-AS-Number on Switch A and Switch B

Switch A:

```
admin@XorPlus# set protocols bgp enable-4byte-as-numbers true
```

Switch B:

```
admin@XorPlus# set protocols bgp enable-4byte-as-numbers true
```

Step 2: Configure BGP-ID and 4-Byte-AS-Number

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "1.0"
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "1.1"
```

Step 3: Configure BGP Peer

Switch A:

```
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "1.1"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "1.0"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 4: Check BGP Peer Status

Switch A:

```
admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/50552 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 65537
Updates Received: 0, Updates Sent: 0
Messages Received: 5, Messages Sent: 5
Time since last received update: n/a
Number of transitions to ESTABLISHED: 7
Time since last entering ESTABLISHED state: 89 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Switch B:

```
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/50552
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 65536
Updates Received: 0, Updates Sent: 0
Messages Received: 6, Messages Sent: 7
Time since last received update: n/a
Number of transitions to ESTABLISHED: 7
Time since last entering ESTABLISHED state: 120 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Step 5: Switch B distribute one BGP route entry to Switch B, then check the BGP route table:

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nextthop	Peer	AS Path
*> 9999::9/128	3003::2	0.0.0.0	?

```
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
    From peer: 0.0.0.0
    Route: Winner
```

```

Origin: INCOMPLETE
AS Path:
Nexthop: 3003::2
Local Preference: 100
admin@XorPlus#

```

Switch A BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

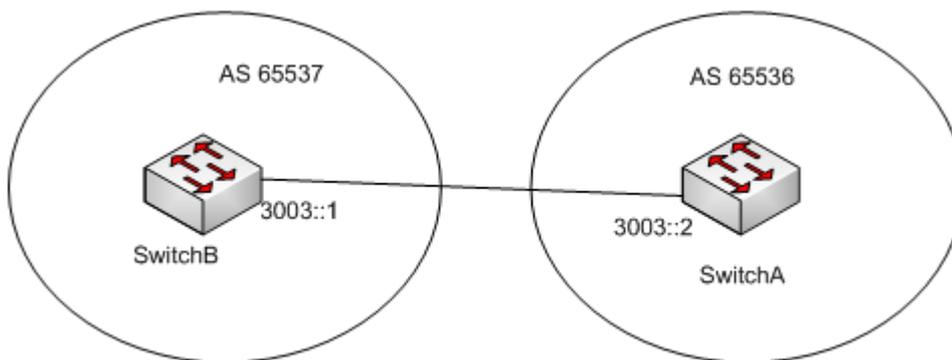
Prefix                Nexthop                Peer                AS Path
-----
*> 9999::9/128        3003::1                6.6.6.6            1.1 ?
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
  From peer: 6.6.6.6
  Route: Winner
  Origin: INCOMPLETE
  AS Path: 1.1
  Nexthop: 3003::1
  Multiple Exit Discriminator: 0
  Local Preference: 100

```

Case 3

Two Peer All Enable 4-byte-as-number

Figure 1-9



Step 1: Enable 4-Byte-AS-Number on Switch A and Switch B

Switch A:

```

admin@XorPlus# set protocols bgp enable-4byte-as-numbers true

```

Switch B:

```

admin@XorPlus# set protocols bgp enable-4byte-as-numbers true

```

Step 2: Configure BGP-ID and 4-Byte-AS-Number

Switch A:

```
admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "65536"
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "65537"
```

Step 3: Configure BGP Peer

Switch A:

```
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "65537"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "65536"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 4: Check BGP Peer Status

Switch A:

```
admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/179 remote 3003::1/52689
  Peer ID: 6.6.6.6
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 65537
  Updates Received: 1, Updates Sent: 0
  Messages Received: 4, Messages Sent: 4
  Time since last received update: 32 seconds
  Number of transitions to ESTABLISHED: 9
  Time since last entering ESTABLISHED state: 32 seconds
  Retry Interval: 120 seconds
  Hold Time: 90 seconds, Keep Alive Time: 30 seconds
  Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
  Minimum AS Origination Interval: 0 seconds
  Minimum Route Advertisement Interval: 0 seconds
```

Switch B:

```
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/52689 remote 3003::2/179
  Peer ID: 9.9.9.9
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 65536
  Updates Received: 0, Updates Sent: 1
  Messages Received: 4, Messages Sent: 4
```

```

Time since last received update: n/a
Number of transitions to ESTABLISHED: 9
Time since last entering ESTABLISHED state: 53 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 5: Switch B distribute one BGP route entry to Switch B, then check the BGP route table:

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::9/128                          3003::2                               0.0.0.0                            ?
admin@XorPlus#
admin@XorPlus#
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
    From peer: 0.0.0.0
    Route: Winner
    Origin: INCOMPLETE
    AS Path:
    Nexthop: 3003::2
    Local Preference: 100

```

Switch A BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::9/128                          3003::1                               6.6.6.6                            1.1 ?
admin@XorPlus# run show bgp routes ipv6 detail
9999::9/128
    From peer: 6.6.6.6
    Route: Winner
    Origin: INCOMPLETE
    AS Path: 1.1
    Nexthop: 3003::1
    Multiple Exit Discriminator: 0
    Local Preference: 100

```

Sources of Routing Updates

Routes can be injected dynamically or statically into BGP. The choice of method depends on the number and stability of routes.

- Injecting Information Dynamically into BGP
- Injecting Information Statically into BGP

Injecting Information Dynamically into BGP

Configuration Command References:

```
admin@XorPlus# set policy policy-statement <Policy-name> term <Term-name> from protocol
<bgp/connected/ospf4/ospf6/rip/static>
```

Note: This command is to specify a policy name and a protocol to be operated.

```
admin@XorPlus# set policy policy-statement <Policy-name> term <Term-name> then <action>
```

Note: This command is to specify an action for the policy-name.

```
admin@XorPlus# set protocols bgp export <Policy-name>
```

Note: This command is to export a special policy which has been defined by the Policy-name.

```
admin@XorPlus# set protocols bgp peer <Peer-IPV6-Address> export <Policy-Name>
```

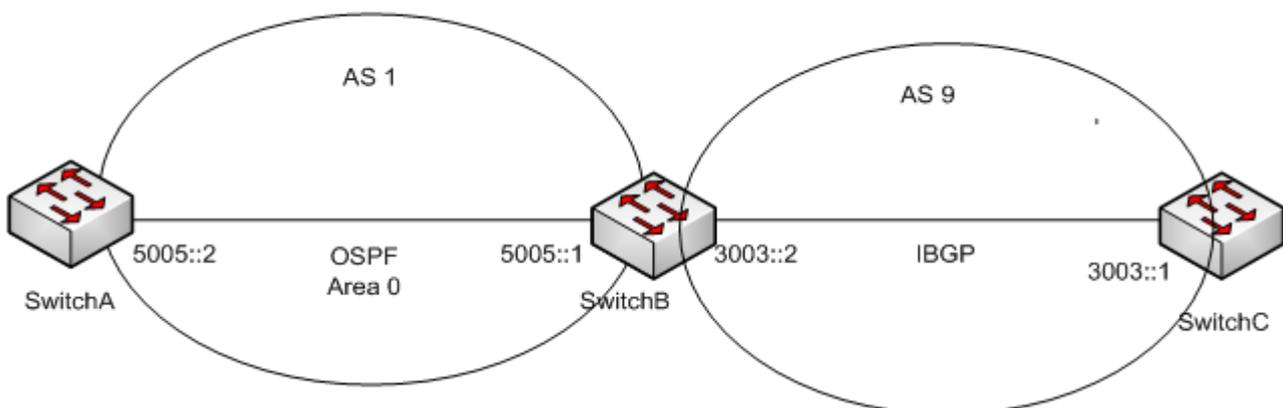
Note: This command is to export the defined policy on a special BGP peer, and the exported BGP route propagates to this special BGP peer.

Inject information dynamically into BGP via route policy, and use export routing policy to inject external route information dynamically into BGP. Reference the section on routing policy for more details about policy.

Configure Example 1:

The following diagram shows that the policy is applied on the global BGP.

Figure 1-10



Step 1: Switch A Switch B enable ospfv3

Switch A:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 1.1.1.1
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::2
```

Switch B:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 9.9.9.9
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::1
```

Step 2: Check ospfv3 neighbor status on Switch B:

Note: OSPFv3 is established.

```
admin@XorPlus# run show ospf6 neighbor
Address                               Interface                               State   Router ID
Pri      Dead
-----
fe80::200:5ff:fe6c:f993              vlan500/vlan500                       Full    1.1.1.1
0          37
```

Step 3: Switch A propagates two ospfv3 routes to Switch B. Check ospf route table on Switch B

Note: Two ospfv3 route entry, 8888::/64 8888:0:0:1::/64, can be seen:

```
admin@XorPlus# run show route table ipv6 unicast ospf
8888::/64                               [ospf(110)/1]
> to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
8888:0:0:1::/64                         [ospf(110)/1]
> to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                               Port
-----
3003::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   connected
5005::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   connected
4001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   te-1/1/46
2001::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   connected
8888::                                     ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93   te-1/1/44
8888:0:0:1::       ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93   te-1/1/44
```

Step 4: Switch B Switch C Enable IBGP**Switch B:**

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "9"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true

```

Step 5: Check BGP peer status

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/50235 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 2, Updates Sent: 0
Messages Received: 6, Messages Sent: 4
Time since last received update: 63 seconds
Number of transitions to ESTABLISHED: 3
Time since last entering ESTABLISHED state: 63 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Switch C:

```

admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/50235
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 0, Updates Sent: 2
Messages Received: 6, Messages Sent: 9
Time since last received update: n/a
Number of transitions to ESTABLISHED: 4
Time since last entering ESTABLISHED state: 109 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 6: Check BGP route table on Switch B. The BGP route table should be NULL.

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
admin@XorPlus#			

Step 7: Configure a policy to Inject ospfv3 route entry into IPV6 BGP on SwitchB

Switch B:

```
admin@XorPlus# set policy policy-statement ospfintobgp term 1 from protocol ospf6
admin@XorPlus# set policy policy-statement ospfintobgp term 1 then accept
```

Step 8: Apply the policy Step 6 defined to BGP

Note: The ospfv3 route entry will inject into BGP route table after applying the policy on BGP

Switch B:

```
admin@XorPlus# set protocols bgp export ospfintobgp
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nextthop	Peer	AS Path
*> 8888::/64	fe80::200:5ff:fe6c:f993	0.0.0.0	?
*> 8888:0:0:1::/64	fe80::200:5ff:fe6c:f993	0.0.0.0	?

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

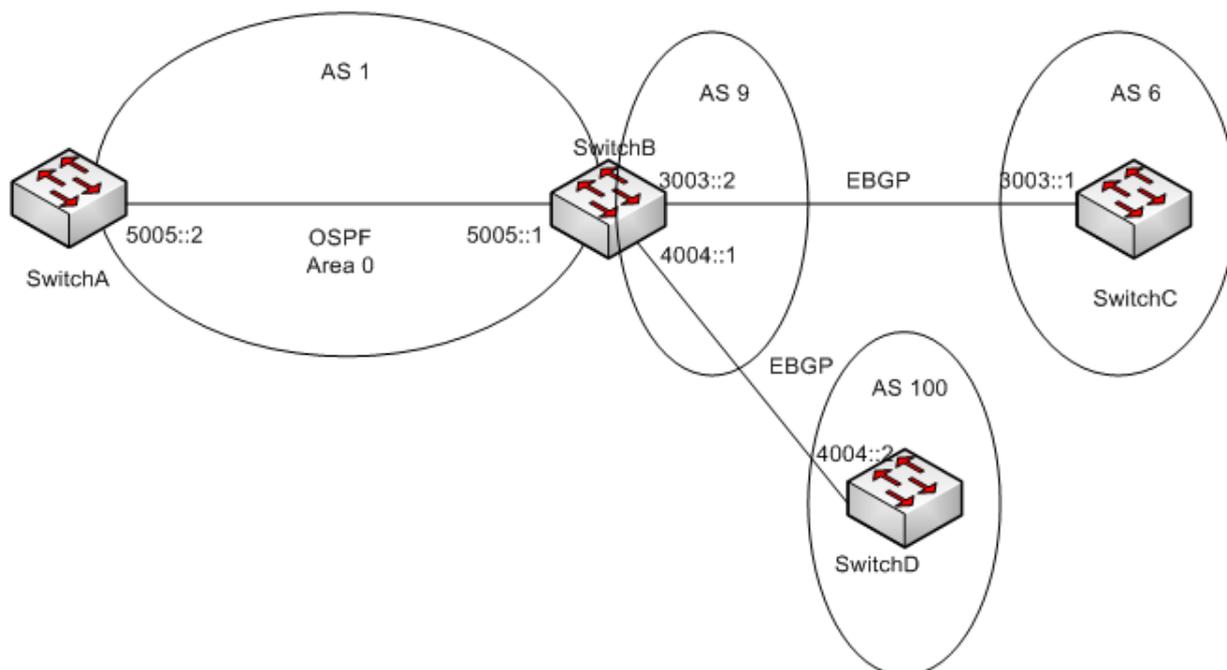
Prefix	Nextthop	Peer	AS Path
*> 8888::/64	3003::2	9.9.9.9	?
*> 8888:0:0:1::/64	3003::2	9.9.9.9	?

```
admin@XorPlus#
```

Configure Example 2:

The following example shows that, with the policy applied on the special peer, BGP route entry will only propagate to this special BGP peer and not to others.

Figure 1-11



Step 1: Switch A Switch B enable ospfv3

Switch A:

```
admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 1.1.1.1
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::2
```

Switch B:

```
admin@XorPlus#set protocols ospf6 instance-id 1
admin@XorPlus#set protocols ospf6 router-id 9.9.9.9
admin@XorPlus#set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::1
```

Step 2: Check ospfv3 neighbor status on Switch B:

Note: OSPFv3 is established.

```
admin@XorPlus# run show ospf6 neighbor
Address                               Interface           State      Router ID
Pri   Dead
-----
fe80::200:5ff:fe6c:f993              vlan500/vlan500    Full      1.1.1.1
0      37
```

Step 3: Switch A propagates two ospfv3 routes to Switch B. Then, check ospf route table on Switch B

Note: Two ospfv3 route entry, 8888::/64 8888:0:0:1::/64, can be seen.

```

admin@XorPlus# run show route table ipv6 unicast ospf
8888::/64      [ospf(110)/1]
                > to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
8888:0:0:1::/64 [ospf(110)/1]
                > to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
admin@XorPlus# run show route forward-route ipv6 all
Destination      NetMask
NextHopMac        Port
-----
3003::           ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
5005::           ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
4001::           ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF te-1/1/46
2001::           ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66 connected
8888::           ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93 te-1/1/44
8888:0:0:1::     ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93 te-1/1/44

```

Step 4: Switch B Switch C Enable EBGP. Switch B Switch D Enable EBGP.

Switch B:

```

admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus#set protocols bgp peer 3003::1 as "6"
admin@XorPlus#set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus#set protocols bgp peer 4004::2 as "100"
admin@XorPlus#set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch C:

```

admin@XorPlus#set protocols bgp bgp-id 6.6.6.6
admin@XorPlus#set protocols bgp local-as "6"
admin@XorPlus#set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus#set protocols bgp peer 3003::2 as "9"
admin@XorPlus#set protocols bgp peer 3003::2 ipv6-unicast true

```

Switch D:

```

admin@XorPlus#set protocols bgp bgp-id 100.100.100.100
admin@XorPlus#set protocols bgp local-as 100
admin@XorPlus#set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus#set protocols bgp peer 4004::1 as 9
admin@XorPlus#set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 5: Check BGP peer status on Switch B

Note: GBP peer all established can be seen.

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/179 remote 3003::1/41512
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 5
Messages Received: 56, Messages Sent: 62
Time since last received update: n/a
Number of transitions to ESTABLISHED: 5
Time since last entering ESTABLISHED state: 1426 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus# run show bgp peers detail 4004::2
Peer 1: local 4004::1/34116 remote 4004::2/179
Peer ID: 100.100.100.100
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 100
Updates Received: 1, Updates Sent: 0
Messages Received: 8, Messages Sent: 7
Time since last received update: 120 seconds
Number of transitions to ESTABLISHED: 2
Time since last entering ESTABLISHED state: 144 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#

```

Step 6: Check BGP route table on Switch B Switch C Switch D. They should all be NULL.

Switch B:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

Switch D:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

Step 7: Configure a policy to Inject ospfv3 route entry into IPV6 BGP on Switch B

Switch B:

```
admin@XorPlus# set policy policy-statement ospfintobgp term 1 from protocol ospf6
admin@XorPlus# set policy policy-statement ospfintobgp term 1 then accept
```

Step 8: Apply the policy on the special BGP peer

Note: BGP route will only propagate to the special BGP peer, not to other peers. If the policy is applied on peer 3003::1, the BGP route will propagate to Switch C but will not propagate to peer 4004::2 on Switch D. Thus, BGP route entry on Switch C can be seen, but the BGP route table on Switch D is still NULL.

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::1 export ospfintobgp
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----
*> 8888::/64	fe80::200:5ff:fe6c:f993	0.0.0.0	?
*> 8888:0:0:1::/64	fe80::200:5ff:fe6c:f993	0.0.0.0	?

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----
*> 8888::/64	3003::2	9.9.9.9	9 ?
*> 8888:0:0:1::/64	3003::2	9.9.9.9	9 ?

Note: The BGP route table just propagates to peer 3003::1, so the BGP route entry from peer 3003::2 can be seen.

Switch D:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

```
admin@XorPlus#
```

Note: The BGP route entry did not propagate to peer 4004::2, so the bgp route table is null on Switch D.

Injecting Information Statically into BGP

Listing prefixes with the network command has the same drawbacks as the dynamic redistribution. If a route that is listed with the network command goes down, BGP will send an update; if the route comes back, BGP will send another update. If this behavior continues, the IGP instability will translate into BGP instabilities. The only way around this is to use a combination of statically defined prefixes in conjunction with the network command. This will ensure that the prefixes will always remain in the IP routing tables and will always be advertised.

Configuration Command References:

```
admin@XorPlus# set protocols bgp network6 <IPV6-Network>
```

Note: This command is to advertise a special IPV6 Network statically.

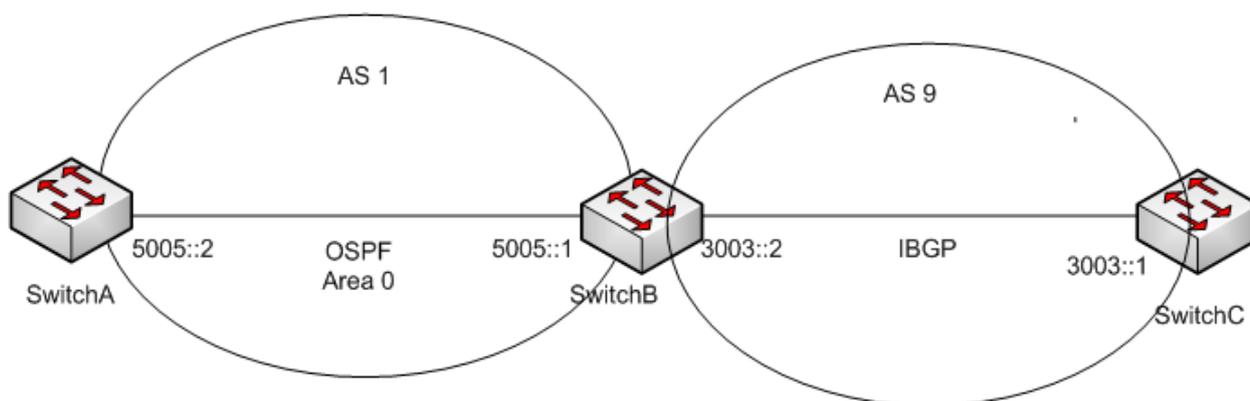


Figure 1-12

Step 1: Switch A Switch B enable ospfv3

Switch A:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 1.1.1.1
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vln500 vif vln500 address
5005::2
```

Switch B:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 9.9.9.9
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vln500 vif vln500 address
5005::1
```

Step 2: Check ospfv3 neighbor status on Switch B

Note: OSPFv3 have been established.

```
admin@XorPlus# run show ospf6 neighbor
Address                               Interface      State      Router ID
Pri      Dead
-----
fe80::200:5ff:fe6c:f993              vlan500/vlan500  Full      1.1.1.1
0          37
```

Step 3: Switch A propagate two ospfv3 route to Switch B. Then, check ospf route table on SwitchB

Note: Two ospfv3 route entries, 8888::/64 8888:0:0:1::/64, can be seen.

```
admin@XorPlus# run show route table ipv6 unicast ospf
8888::/64      [ospf(110)/1]
                > to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
8888:0:0:1::/64 [ospf(110)/1]
                > to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac      Port
-----
3003::          ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66  connected
5005::          ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66  connected
4001::          ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF te-1/1/46
2001::          ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66  connected
8888::          ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93 te-1/1/44
8888:0:0:1::    ffff:ffff:ffff:ffff::
00:00:05:6C:F9:93 te-1/1/44
```

Step 4: Switch B Switch C Enable EBGP. Switch B Switch D Enable EBGP.

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 5: Check BGP peer status on Switch B

Note: BGP peer is established.

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/179 remote 3003::1/41512
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 5
Messages Received: 56, Messages Sent: 62
Time since last received update: n/a
Number of transitions to ESTABLISHED: 5
Time since last entering ESTABLISHED state: 1426 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 6: Check BPG route table on Switch B Switch C Switch D. They should all be NULL.

Switch B:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----

Step 7: Injecting a special network into BGP via “Network” command

Note: We inject network 8888::/64 into BGP via network command, and the network 8888::/64 is reachable, as the entry 8888::/64 is in ospfv3 route table.

```

admin@XorPlus# set protocols bgp network6 8888::/64

```

Step 8: Check BGP route table on Switch B Switch C

Switch B:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	NextHop	Peer	AS Path
-----	-----	-----	-----
*> 8888::/64	::	0.0.0.0	i

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 8888::/64	3003::2	9.9.9.9	9 i

Step 9: Switch A withdrawn the route entry 8888::/64

Since the network 8888::/64 is not reachable, the network will not install into the BGP route table.

Switch B ospfv3 route table:

```
admin@XorPlus# run show route table ipv6 unicast ospf
admin@XorPlus#
```

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
admin@XorPlus#			

BGP Attributes

In this section will discuss how the different BGP attributes are used.

- The NEXT_HOP Attribute
- The AS_PATH Attribute
- The LOCAL_PREF Attribute
- The MULTI_EXIT_DISC Attribute
- The COMMUNITY Attribute

The NEXT_HOP Attribute

Configuration Example 1

Case 1 :

The NEXT_HOP will be the NEXT_HOP of the external Neighbor.

As the following example illustrates, Switch A distributes 7777::/64 to Switch B. As seen, the next_hop is 5005::02 on both Switch B and Switch C. It will not be active, as Switch C does not know how to reach 5005::2. Thus, user must instruct Switch C on how to reach next_hop, 5005::2 via IGP.

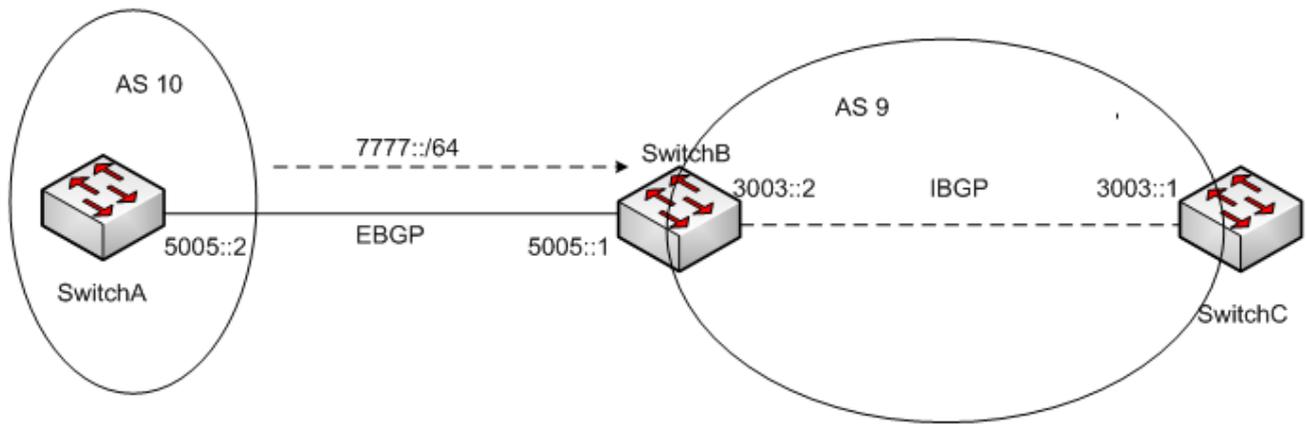


Figure 1-13

Step 1: Switch A and Switch B Enable EBGP. Switch B and Switch C Enable IGP.

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "9"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 2: Switch A propagate a BGP route entry 7777::/64 to Switch B

Note: Check the BGP route table on Switch B. The next_hop should be 5005::2. And the next_hop also is 5005::2, but it's not reachable on Switch C, so the BGP route entry 7777::/64 is inactive.

Switch B:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
--------	---------	------	---------

```

-----
*> 7777::/64                               5005::2                               33.33.33.33                               10 i
-----

```

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

```

-----
Prefix                               Nexthop                               Peer                               AS Path
-----
* 7777::/64                           5005::2                               9.9.9.9                           10 i
-----

```

Step 3: Configure a static route to the destination 5005::2/64, and specify the next_hop 3003::1/64 on Switch C**Switch C:**

```

admin@XorPlus# set protocols static route 5005::/64 next-hop 3003::2

```

```

admin@XorPlus# run show route forward-route ipv6 all

```

```

Destination                               NetMask
NextHopMac                               Port
-----
2001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   connected
3003::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   connected
5005::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   te-1/1/46
7777::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   te-1/1/46

```

Step 4: Check BGP route table on Switch C, the BGP route entry 7777::/64 is already active as the next_hop 5005::2 was reachable and it is installed into route table.

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

```

-----
Prefix                               Nexthop                               Peer                               AS Path
-----
*> 7777::/64                           5005::2                               9.9.9.9                           10 i
-----

```

```

admin@XorPlus#

```

```

admin@XorPlus# run show route forward-route ipv6 all

```

```

Destination                               NetMask
NextHopMac                               Port
-----
2001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   connected
3003::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF   connected
5005::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   te-1/1/46
7777::                                     ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66   te-1/1/46

```

Configuration Example 2

Case 2:

Use the `set protocols bgp peer <Peer-IPV6-Address> next-hop-self true` command to force the router to advertise itself.

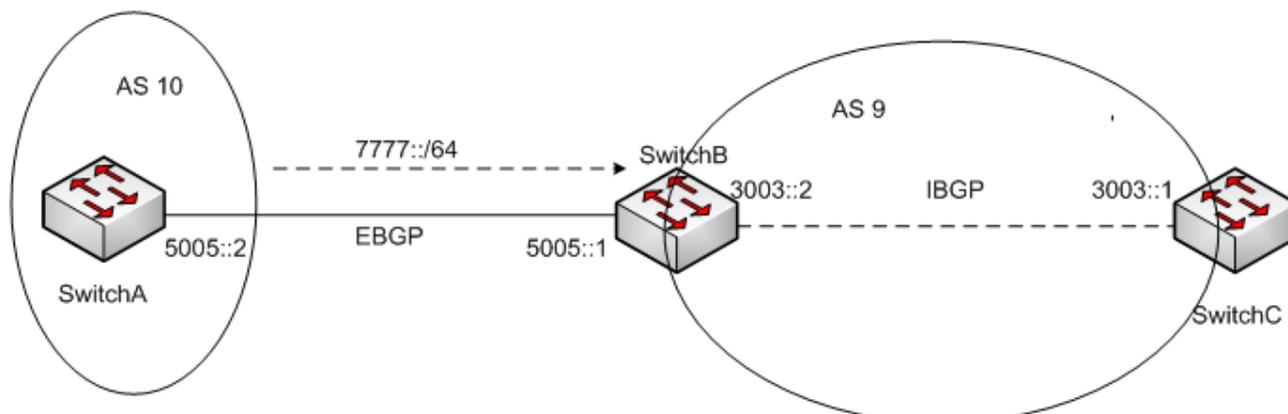


Figure 1-14

Step 1: Switch A Switch B Enable EBGP, Switch B Switch C Enable IGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "9"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 2: Switch A propagate a BGP route entry 7777::/64 to Switch B

Note: Check the BGP route table on Switch B. The next_hop should be 5005::2.

And the next_hop also is 5005::2 and it's not reachable on SwitchC, so the bgp route entry 7777::/64 is inactive.

Switch B:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 7777::/64	5005::2	33.33.33.33	10 i

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
* 7777::/64	5005::2	9.9.9.9	10 i

Step 3: Enable Next-hop-self on Switch B

Note: After enabling next-hop-self on Switch B, the next-hop will be forced to be modified by itself using the IPv6 address.

Switch B:

```
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 7777::/64	5005::2	33.33.33.33	10 i

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 7777::/64	3003::2	9.9.9.9	10 i

Note: We can see that the next-hop has been set as 3003::2.

Configuration Example 3

Case 3:

The NEXT_HOP always is the peer IPv6 address who advertised it.

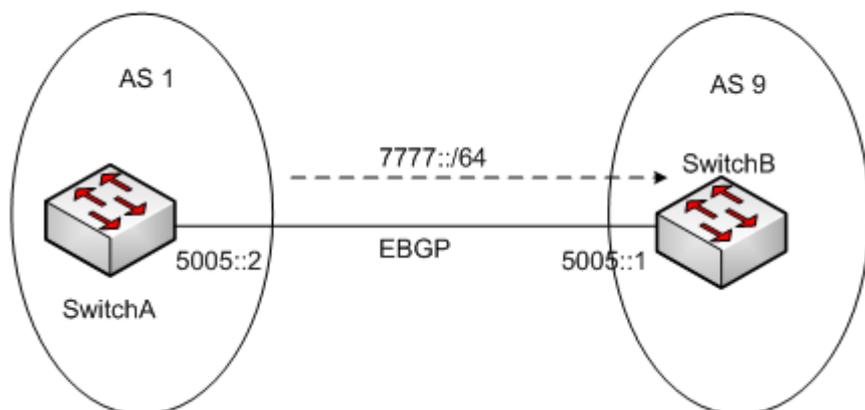


Figure 1-15

Step 1: Switch A Switch B Enable EBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Step 2: Switch A propagate a BGP route entry 7777::/64 to Switch B

Note: Check the BGP route table on Switch B, The next_hop should be 5005::2.

Switch B:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 7777::/64	5005::2	33.33.33.33	10 i

The AS_PATH Attribute

AS_PATH is a list of AS-numbers that the BGP route passed. Router will elect the shortest AS_PATH as the best path when other attributes are the same.

If using private AS-numbers (64512-65534), user can remove the private AS-number from the AS_PATH to prevent the leakage of private AS numbers into the Internet.

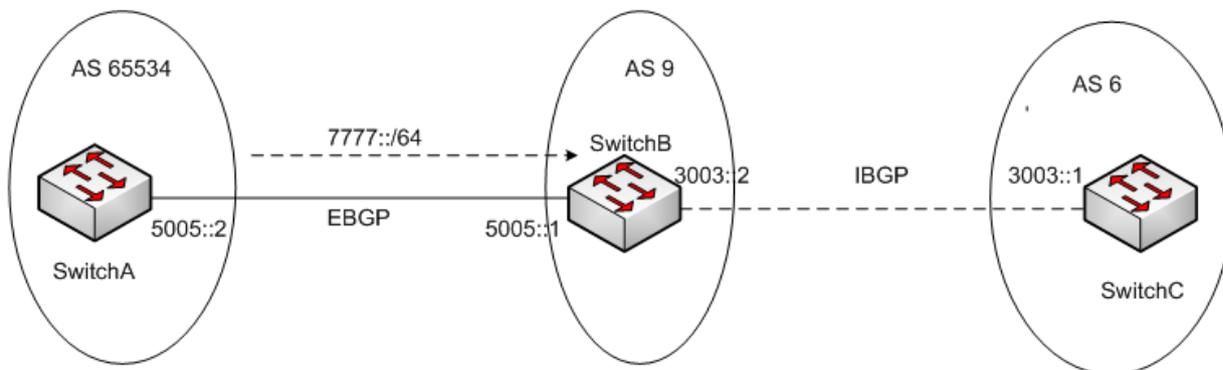


Figure 1-16

Step 1: Switch A and Switch B Enable EBGP. Switch B and Switch C Enable EBGP. Switch A uses a private AS-number 65534

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# et protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 65534
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "65534"
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "65534"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 2: Check BGP peer status on Switch B

Switch B:

```

admin@XorPlus# run show bgp peers detail 3003::1
Peer 1: local 3003::2/50985 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 1
Messages Received: 4, Messages Sent: 5
Time since last received update: n/a
Number of transitions to ESTABLISHED: 8
Time since last entering ESTABLISHED state: 69 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus# run show bgp peers detail 5005::2
Peer 1: local 5005::1/36229 remote 5005::2/179
Peer ID: 33.33.33.33
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 10
Updates Received: 1, Updates Sent: 0
Messages Received: 9, Messages Sent: 8
Time since last received update: 156 seconds
Number of transitions to ESTABLISHED: 3
Time since last entering ESTABLISHED state: 156 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 3: Switch A propagate a BGP route entry 7777::/64

Note :The AS-path should be 65534 on Switch B, and the AS-path should be 9 65534 on Switch C.

Switch B:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 7777::/64                           5005::2                               33.33.33.33                       10 i

```

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 7777::/64                           3003::2                               9.9.9.9                           65534 10 i

```

Step 4: Remove private AS-number from AS_PATH

Note: The private AS-number should be removed from the AS-path after enabled public-as-only. It only removes the private AS-number when the private AS-number is its own.

Switch B:

```

admin@XorPlus# set protocols bgp peer 3003::1 public-as-only true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 7777::/64                           5005::2                               33.33.33.33                       10 i
admin@XorPlus#

```

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 7777::/64                           3003::2                               9.9.9.9                           10 i
admin@XorPlus#

```

Commands References

```
admin@XorPlus# set protocols bgp local-preference <0- 4294967295>
```

Note: This command is to set the value of local-preference. It also affects the BGP decision process. If multiple paths for the same prefix are available, the path with the larger local preference value is preferred. LOCAL_PREF is an AS-wide attribute at the highest level of the BGP decision process. It is considered before the AS path length. A longer path with a larger local preference is preferred over a shorter path with a smaller local preference.

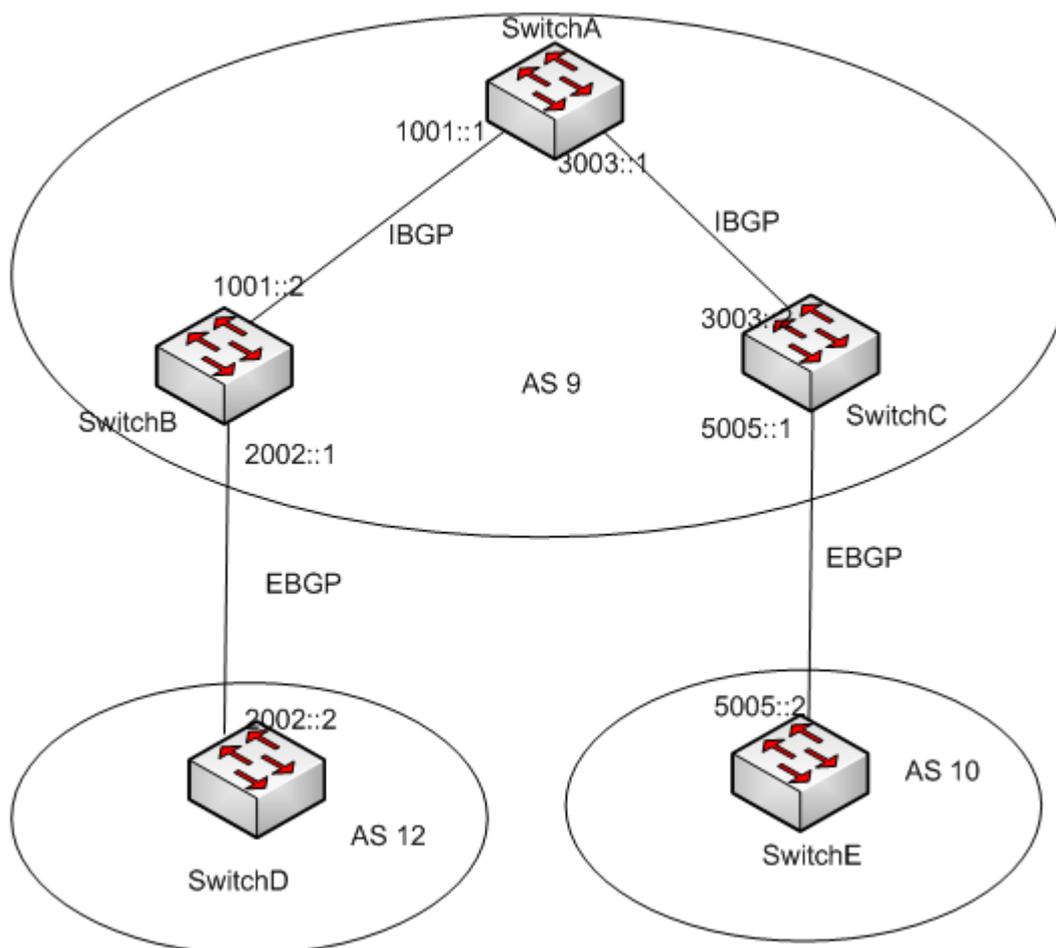


Figure 1-17

Step 1: Configure BGP as TOPO displayed

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "9"
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as 9
admin@XorPlus# set protocols bgp peer 1001::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip "2002::1"
admin@XorPlus# set protocols bgp peer 2002::2 as "12"
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as 9
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true

```

Switch D:

```

admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 12
admin@XorPlus# set protocols bgp peer 2002::1 local-ip "2002::2"
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true

```

Switch E:

```

admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true

```

Step 2: Switch D and Switch E propagate BGP route entry 9999::/64 to Switch B and Switch C

Note: check the BGP route table on Switch A.

Switch B:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	2002::2	44.44.44.44	12 i

Switch C:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	5005::2	33.33.33.33	10 i

Switch A:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
--------	---------	------	---------

```

-----
* 9999::/64          1001::2          26.26.26.26      12 i
*> 9999::/64        3003::2          9.9.9.9          10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 10
    Nexthop: 3003::2
    Local Preference: 100
9999::/64
    From peer: 26.26.26.26
    Route: Not Used
    Origin: IGP
    AS Path: 12
    Nexthop: 1001::2
    Local Preference: 100
admin@XorPlus#

```

Note: The Local-Preference all is 100 the two BGP route entry.

Step 3: Modify the Local-Preference value on Switch B

Switch B:

```
admin@XorPlus# set protocols bgp local-preference 200
```

Switch A:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix          Nexthop          Peer          AS Path
-----
* 9999::/64      3003::2          9.9.9.9       10 i
*> 9999::/64     1001::2          26.26.26.26  12 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Not Used
    Origin: IGP
    AS Path: 10
    Nexthop: 3003::2
    Local Preference: 100
9999::/64
    From peer: 26.26.26.26
    Route: Winner
    Origin: IGP
    AS Path: 12
    Nexthop: 1001::2
    Local Preference: 200
admin@XorPlus#

```

Note: The best route is coming from Switch B, as the BGP route entry from Switch B has a larger Local-Preference value. It will select the BGP route entry with smaller Local-Preference value if other attributes all have the same priority.

Step 4: Modify the Local-Preference larger than Switch B on Switch C

Switch C:

```
admin@XorPlus# set protocols bgp local-preference 300
```

Switch A:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	1001::2	26.26.26.26	12 i
*> 9999::/64	3003::2	9.9.9.9	10 i

```
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 3003::2
  Local Preference: 300
9999::/64
  From peer: 26.26.26.26
  Route: Not Used
  Origin: IGP
  AS Path: 12
  Nexthop: 1001::2
  Local Preference: 200
admin@XorPlus#
```

Note: The BGP speaker selects the BGP route entry with larger Local-Preference.

The MULTI_EXIT_DISC Attribute

This section demonstrates how MED can be used by one AS to influence routing decisions of another AS.

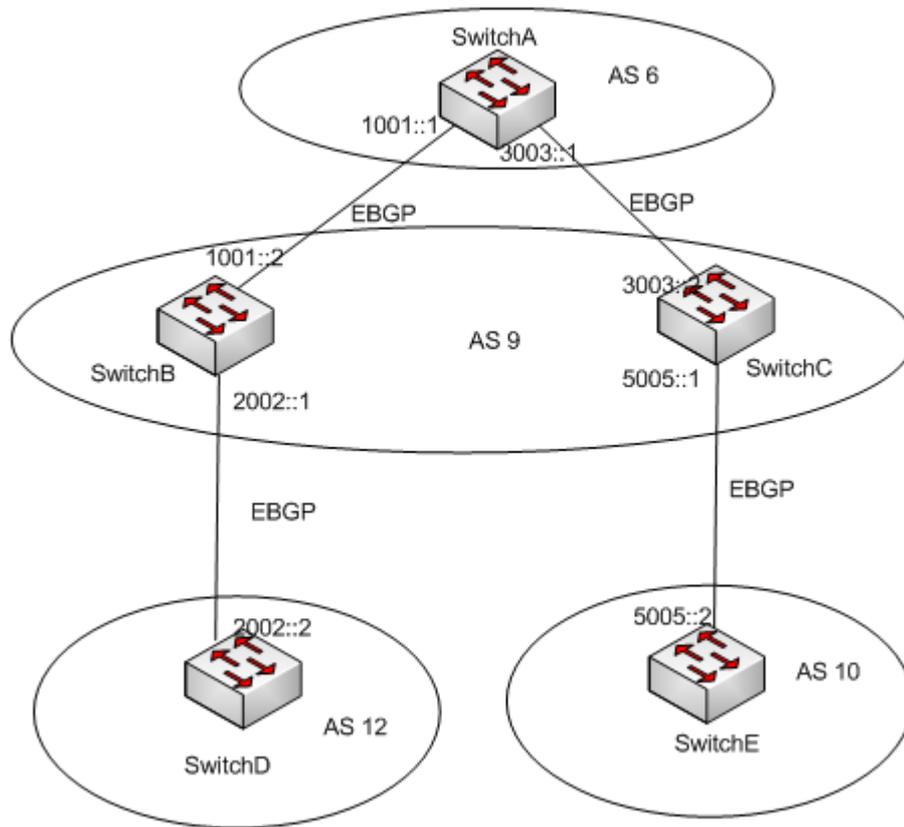


Figure 1-18

Step 1: Configure BGP as TOPO displayed

Switch A:

```
admin@XorPlus#set protocols bgp bgp-id 6.6.6.6
admin@XorPlus#set protocols bgp local-as "6"
admin@XorPlus#set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus#set protocols bgp peer 1001::2 as "9"
admin@XorPlus#set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus#set protocols bgp peer 3003::2 as "9"
admin@XorPlus#set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus#set protocols bgp bgp-id 26.26.26.26
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus#set protocols bgp peer 1001::1 as "6"
admin@XorPlus#set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 2002::2 local-ip "2002::1"
admin@XorPlus#set protocols bgp peer 2002::2 as "12"
admin@XorPlus#set protocols bgp peer 2002::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus#set protocols bgp bgp-id 9.9.9.9
admin@XorPlus#set protocols bgp local-as "9"
admin@XorPlus#set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus#set protocols bgp peer 3003::1 as "6"
```

```
admin@XorPlus#set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus#set protocols bgp peer 5005::2 as "10"
admin@XorPlus#set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch D:

```
admin@XorPlus#set protocols bgp bgp-id 44.44.44.44
admin@XorPlus#set protocols bgp local-as 12
admin@XorPlus#set protocols bgp peer 2002::1 local-ip "2002::2"
admin@XorPlus#set protocols bgp peer 2002::1 as 9
admin@XorPlus#set protocols bgp peer 2002::1 ipv6-unicast true
```

Switch E:

```
admin@XorPlus#set protocols bgp bgp-id 33.33.33.33
admin@XorPlus#set protocols bgp local-as 10
admin@XorPlus#set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus#set protocols bgp peer 5005::1 as 9
admin@XorPlus#set protocols bgp peer 5005::1 ipv6-unicast true
```

Step 2: Switch D and Switch E propagate BGP route entry 9999::/64 to Switch B and Switch C

Note: check the BGP route table on Switch A.

Switch B:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	2002::2	44.44.44.44	12 i

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	5005::2	33.33.33.33	10 i

Switch A:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	1001::2	26.26.26.26	9 12 i
*> 9999::/64	3003::2	9.9.9.9	9 10 i

```
admin@XorPlus#
```

```
admin@XorPlus# run show bgp routes ipv6 de
```

```
Possible completions:
```

```
<IPNet>
```

```
Print BGP IPv6 routes of specified prefix
```

```

detail                               Print detailed BGP IPv6 routes
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 9 10
    Nexthop: 3003::2
    Multiple Exit Discriminator: 0
    Local Preference: 100
9999::/64
    From peer: 26.26.26.26
    Route: Not Used
    Origin: IGP
    AS Path: 9 12
    Nexthop: 1001::2
    Multiple Exit Discriminator: 0
    Local Preference: 100

```

Note: The MED all is 0. The two BGP route entry.

Step 3: Modify the MED value on Switch C

Switch C:

```
admin@XorPlus# set protocols bgp med 100
```

Switch A:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                Nexthop                Peer                AS Path
  -----                -
* 9999::/64             3003::2                9.9.9.9             9 10 i
*> 9999::/64            1001::2                26.26.26.26         9 12 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 26.26.26.26
    Route: Winner
    Origin: IGP
    AS Path: 9 12
    Nexthop: 1001::2
    Multiple Exit Discriminator: 0
    Local Preference: 100
9999::/64
    From peer: 9.9.9.9
    Route: Not Used
    Origin: IGP
    AS Path: 9 10
    Nexthop: 3003::2
    Multiple Exit Discriminator: 100
    Local Preference: 100
admin@XorPlus#

```

Note: The best route is coming from Switch B, as the BGP route entry from Switch B has a smaller MED value. It will select the BGP route entry with smaller MED value if other attributes all have the same priority.

Step 4: Modify the MED larger than Switch C

Switch B:

```
admin@XorPlus# set protocols bgp med 200
```

Switch A:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	1001::2	26.26.26.26	9 12 i
*> 9999::/64	3003::2	9.9.9.9	9 10 i

```
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 3003::2
  Multiple Exit Discriminator: 100
  Local Preference: 100
9999::/64
  From peer: 26.26.26.26
  Route: Not Used
  Origin: IGP
  AS Path: 9 12
  Nexthop: 1001::2
  Multiple Exit Discriminator: 200
  Local Preference: 100
```

Note: The BGP speaker selects the BGP route entry with the smaller MED.

Command References

```
admin@XorPlus# set protocols bgp peer <Peer-ipv6-address> advertise community disable
<true/false>
```

Note: This command is to enable the BGP speaker to advertise well known community attribute NO_EXPORT (0xFFFFFFFF01) NO_ADVERTISE (0xFFFFFFFF02) NO_EXPORT_SUBCONFED (0xFFFFFFFF03).

```
admin@XorPlus# set protocols bgp peer <Peer-ipv6-address> advertise community-ext disable
<true/false>
```

Note: This command is to enable BGP speakers advertising unknown community attributes.

A community is a group of destinations which share some common property. Each autonomous system administrator may define which communities a destination belongs to. By default, all destinations belong to the general Internet community.

The COMMUNITIES path attribute is an optional transitive attribute of variable length. The attribute consists of a set of four octet values, each of which specify a community. All routes with this attribute belong to the communities listed in the attribute.

The COMMUNITIES attribute has Type Code 8.

Communities are treated as 32 bit values, however for administrative assignment, the following presumptions may be made:

The community attribute values ranging from 0x00000000 through 0x0000FFFF, and 0xFFFF0000 through 0xFFFFFFFF, are hereby reserved. The rest of the community attribute values shall be encoded using an autonomous system number in the first two octets. The semantics of the final two octets may be defined by the autonomous system (e.g. AS690 may define research, educational and commercial community values that may be used for policy routing as defined by the operators of that AS, using community attribute values 0x02B20000 through 0x02B2FFFF).

A BGP speaker may use this attribute to control which routing information it accepts, prefers or distributes to other neighbors.

A BGP speaker receiving a route that does not have the COMMUNITIES path attribute may append this attribute to the route when propagating it to its peers.

A BGP speaker receiving a route with the COMMUNITIES path attribute may modify this attribute according to the local policy.

Aggregation

If a range of routes is to be aggregated and the resultant aggregate attribute section does not carry the ATOMIC_AGGREGATE attribute, then the resulting aggregate should have a COMMUNITIES path attribute which contains all communities from all of the aggregated routes.

Well-known Communities

The following communities have global significance and their operations shall be implemented in any community-attribute-aware BGP speaker.

NO_EXPORT (0xFFFFFFFF01):

All routes received carrying a communities attribute containing this value MUST NOT be advertised outside a BGP confederation boundary (a stand-alone autonomous system that is not part of a confederation should be considered a confederation itself).

NO_ADVERTISE (0xFFFFFFFF02)

All routes received carrying a communities attribute containing this value MUST NOT be advertised to other BGP peers.

NO_EXPORT_SUBCONFED (0xFFFFFFFF03)

All routes received carrying a communities attribute containing this value MUST NOT be advertised to external BGP peers (this includes peers in other members autonomous systems inside a BGP confederation).

Because communities are not propagated to internal or external BGP neighbors by default, the command **set protocols bgp peer <peer-ipv6-address> advertise community-ext disable true** and **set protocols bgp peer <peer-ipv6-address> advertise community disable true** in order for the assigned community to be sent out.

Configuration Example 1

The following example is to show you that, 7777::/64 will not advertise to EBGP peer, 8888::/64 will not advertise to any BGP peer (IBGP,EBGP), and 9999::/64 will not advertise to EBGP BGP peer (including confederation EBGP peer).

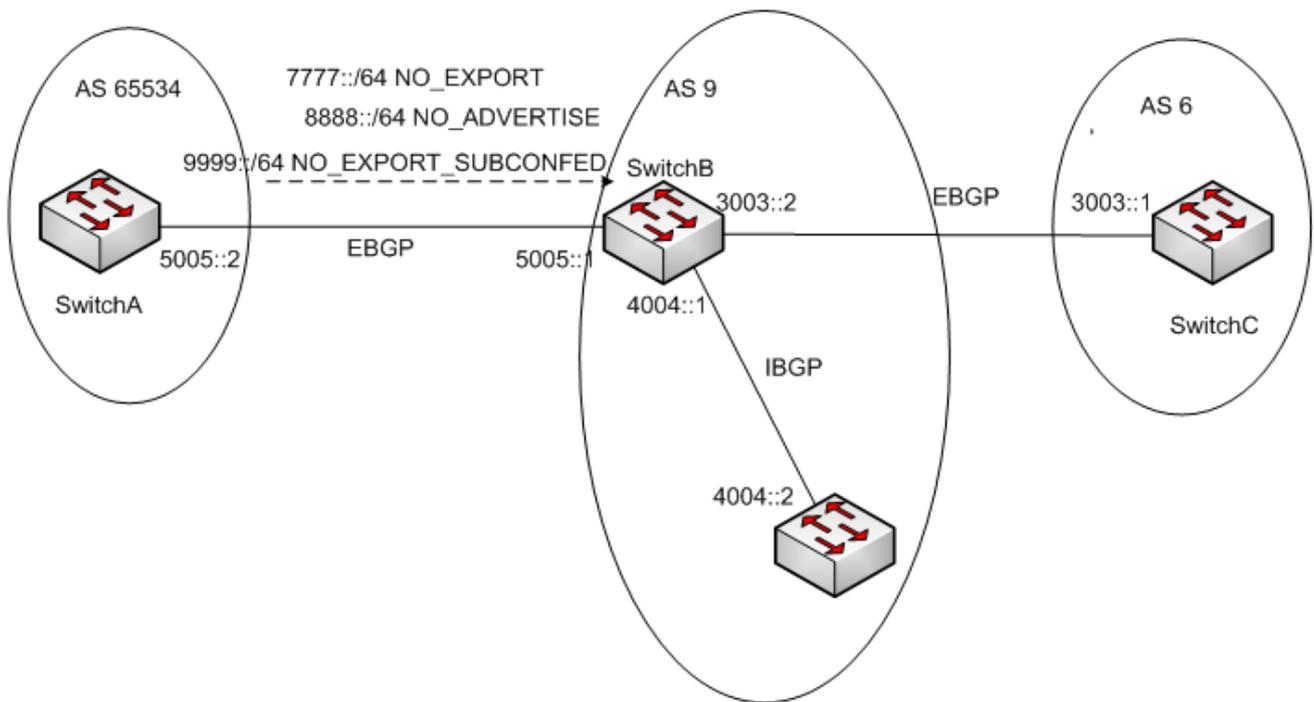


Figure 1-19

Step 1: Switch A Switch B Enable EBGP, Switch B Switch C Enable EBGP, Switch B Switch D Enable IBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 3003::1 local-ip 3003::2
admin@XorPlus# set protocols bgp peer 3003::1 as 6
admin@XorPlus# set protocols bgp peer 3003::1 public-as-only true
admin@XorPlus# set protocols bgp peer 3003::1 advertise community disable false
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip 4004::1
admin@XorPlus# set protocols bgp peer 4004::2 as 9
admin@XorPlus# set protocols bgp peer 4004::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip 5005::1
admin@XorPlus# set protocols bgp peer 5005::2 as 10
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as 6
admin@XorPlus# set protocols bgp peer 3003::2 local-ip 3003::1
admin@XorPlus# set protocols bgp peer 3003::2 as 9
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true

```

Switch D:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 4004::1 local-ip 4004::2
admin@XorPlus# set protocols bgp peer 4004::1 as 9
admin@XorPlus# set protocols bgp peer 4004::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: Switch A propagates three BGP route entries to Switch B

Note: Switch A propagates three BGP route entries 7777::/64 with community NO_EXPORT, it will not advertise to any EBGP peer, 8888::/64 with community NO_ADVERTISE, 9999::/64 will not advertise to any BGP peer, with community NO_EXPORT_SUBCONFED, it will not advertise to any EBGP peer including confederation EBGP peer. So Switch C will receive no BGP route entry, as Switch B and Switch C is an EBGP peer, but Switch D will receive 7777::/64 and 9999::/64, because, Switch B and Switch D are IBGP peers, as 8888::/64 with community NO_ADVERTISE, it will not advertise to any BGP peer, so Switch D just can't receive 8888::/64.

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                                   AS Path
  -----                               -
*> 7777::/64                            5005::2                               33.33.33.33                           10 i
*> 8888::/64                            5005::2                               33.33.33.33                           10 i
*> 9999::/64                            5005::2                               33.33.33.33                           10 i
admin@XorPlus# run show bgp routes ipv6 detail
7777::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 300
  Community: 0xffffffff01[NO_EXPORT]
8888::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 300
  Community: 0xffffffff02[NO_ADVERTISE]
9999::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10

```

```

Nexthop: 5005::2
Local Preference: 300
Community: 0xffffffff03[NO_EXPORT_SUBCONFED]
admin@XorPlus#

```

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path

```
admin@XorPlus#
```

Switch D BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path

*> 7777::/64	4004::1	9.9.9.9	10 i
*> 9999::/64	4004::1	9.9.9.9	10 i

```
admin@XorPlus# run show bgp routes ipv6 detail
```

```
7777::/64
```

```

From peer: 9.9.9.9
Route: Winner
Origin: IGP
AS Path: 9 10
Nexthop: 4004::1
Local Preference: 300

```

```
9999::/64
```

```

From peer: 9.9.9.9
Route: Winner
Origin: IGP
AS Path: 9 10
Nexthop: 4004::1
Local Preference: 300

```

```
admin@XorPlus#
```

Step 3: Enable advertise community on Switch B

Note: Community value will not advertise to it's BGP peer by default, but we can Enable BGP speakers to advertise community to it's BGP peer, the command "**set protocols bgp peer <peer-ipv6-address> advertise community-ext disable true**" will advertise unknown community to it's BGP peer and the command "**set protocols bgp peer <peer-ipv6-address> advertise community disable true**" will advertise a well-known community to it's BGP peer.

Switch B:

```
admin@XorPlus# set protocols bgp peer 4004::2 advertise community disable false
```

Note: The community will be sent out after Enable advertise community.

Step 4: Check the community value of the BGP route entry on Switch D

Switch D:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                               Nexthop                               Peer                                   AS Path
  -----                               -
*> 7777::/64                           4004::1                               9.9.9.9                               10 i
*> 9999::/64                           4004::1                               9.9.9.9                               10 i
admin@XorPlus# run show bgp routes ipv6 detail
7777::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 4004::1
  Local Preference: 300
  Community: 0xffffffff01[NO_EXPORT]
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 4004::1
  Local Preference: 300
  Community: 0xffffffff03[NO_EXPORT_SUBCONFED]
admin@XorPlus#

```

Note: We can see the community after Step 3

Configuration Example 2

This example is to show you that the BGP route entry with Community NO_EXPORT_SUBCONFED will not advertise to EBGP peer, including Confederation EBGP peer. For more detailed information about Confederation, please reference the Confederation section.

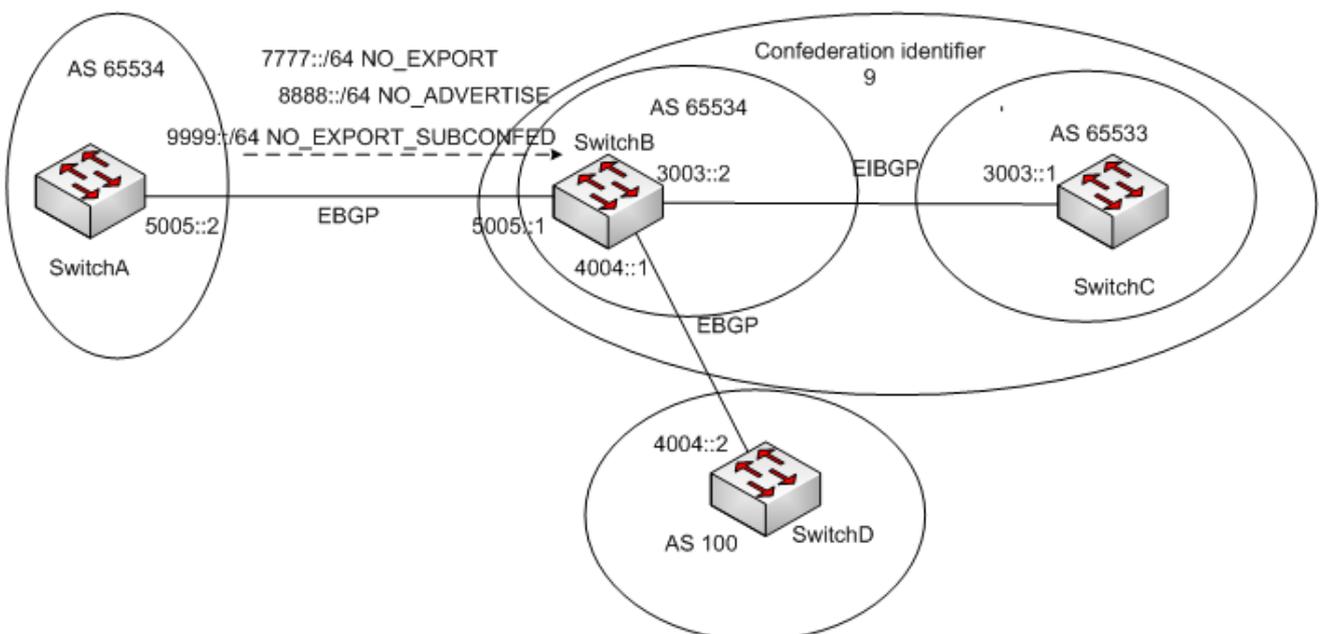


Figure 1-20

Step 1: Switch A Switch B Enable EBGP, Switch B Switch C Enable Confederation, Switch B Switch D Enable EBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as 65534
admin@XorPlus# set protocols bgp confederation identifier 9
admin@XorPlus# set protocols bgp peer 3003::1 local-ip 3003::2
admin@XorPlus# set protocols bgp peer 3003::1 as 65533
admin@XorPlus# set protocols bgp peer 3003::1 confederation-member true
admin@XorPlus# set protocols bgp peer 3003::1 public-as-only true
admin@XorPlus# set protocols bgp peer 3003::1 advertise community disable false
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip 4004::1
admin@XorPlus# set protocols bgp peer 4004::2 as 100
admin@XorPlus# set protocols bgp peer 4004::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::2 advertise community disable false
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip 5005::1
admin@XorPlus# set protocols bgp peer 5005::2 as 10
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as 65533
admin@XorPlus# set protocols bgp peer 3003::2 local-ip 3003::1
admin@XorPlus# set protocols bgp peer 3003::2 as 65534
admin@XorPlus# set protocols bgp peer 3003::2 confederation-member true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip 4004::2
admin@XorPlus# set protocols bgp peer 4004::1 as 9
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true
```

Step 2: Switch A propagate three BGP routes to Switch B

Note: Switch A propagates three BGP route entries to Switch B :7777::/64 with community NO_EXPORT, 8888::/64 with community NO_ADVERTISE, and 9999::/64 with community NO_EXPORT_SUBCONFED. 7777::/64 will not advertise to EBGP Peer, so Switch D will not receive this BGP route entry. However, Switch C will receive this entry, 8888::/64 will not advertise to any BGP Peer, so Switch C Switch D all will not receive this entry, 9999::/64 will not advertise to EBGP Peer, including Confederation EBGP Peer, so Switch C Switch D all will not receive this entry.

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 7777::/64            5005::2                33.33.33.33        10 i
*> 8888::/64            5005::2                33.33.33.33        10 i
*> 9999::/64            5005::2                33.33.33.33        10 i
admin@XorPlus# run show bgp routes ipv6 detail
7777::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 300
  Community: 0xffffffff01[NO_EXPORT]
8888::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 300
  Community: 0xffffffff02[NO_ADVERTISE]
9999::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 300
  Community: 0xffffffff03[NO_EXPORT_SUBCONFED]
admin@XorPlus#

```

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
* 7777::/64            5005::2                9.9.9.9            10 i
admin@XorPlus# run show bgp routes ipv6 detail
7777::/64
  From peer: 9.9.9.9
  Route: Not Used
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Local Preference: 100
  Community: 0xffffffff01[NO_EXPORT]
admin@XorPlus#

SwitchD bgp route table:
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
admin@XorPlus#

```

Configure Example 3:

This example is to show you how to set community attribute values, via policy.

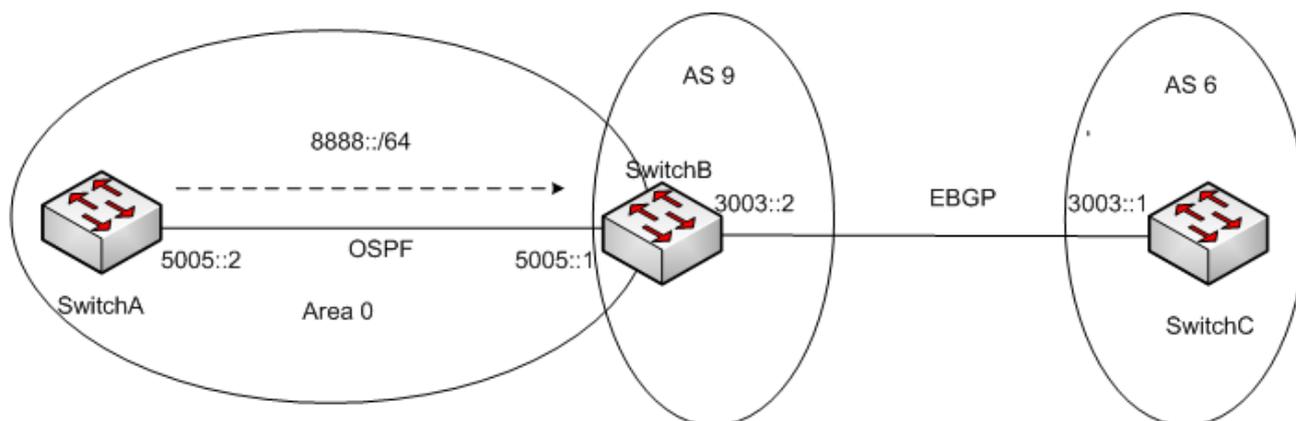


Figure 1-21

Step 1: Switch A Switch B Establish OSPFV3

Note: Switch A Switch B Establish ospv3 and SwitchA propagate one ospfv3 route 8888::/64 to Switch B.

Switch A:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 1.1.1.1
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::2
```

Switch B:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 9.9.9.9
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan500 vif vlan500 address
5005::1
```

Step 2: Switch B Switch C Establish EBGP

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
```

```
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 3: Check ospf route table on Switch B:

```
admin@XorPlus# run show route table ipv6 unicast ospf
8888::/64 [ospf(110)/1]
> to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
8888:0:0:1::/64 [ospf(110)/1]
> to fe80::200:5ff:fe6c:f993 via vlan500/vlan500
```

Step 4: Configure a policy to export ospfv3 into bgp and set community 9:6 to this bgp route entry

Switch B:

```
admin@XorPlus# set policy policy-statement ospfintobgp term 1 from protocol "ospf6"
admin@XorPlus# set policy policy-statement ospfintobgp term 1 to origin 2
admin@XorPlus# set policy policy-statement ospfintobgp term 1 then community "9:6"
```

Step 5: Apply the policy to BGP on Switch B

Note: The unknown community will not advertise to it's BGP peer, so you must enable "advertise community-ext".

Switch B:

```
admin@XorPlus# set protocols bgp export ospfintobgp
admin@XorPlus# set protocols bgp peer 3003::1 advertise community-ext disable false
```

Step 6: Check the community of the BGP route table on Switch C, it should be 9:6

Switch C:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 8888::/64            3003::2                9.9.9.9             9 ?
admin@XorPlus# run show bgp routes ipv6 detail
8888::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: INCOMPLETE
  AS Path: 9
  Nexthop: 3003::2
  Multiple Exit Discriminator: 0
  Local Preference: 100
  Community: 0x90006[9:6]
admin@XorPlus#
```

Command References

```
admin@XorPlus# set protocols bgp aggregate network6 <IPv6-Network>
```

Note: This command is used to set IPv6 route aggregation. Routes with different MED values should not be aggregated.

```
admin@XorPlus# set protocols bgp aggregate network6 <IPv6-Network> suppress-detail  
<true/false>
```

Note: This command is used to enable or disable suppress-detail BGP route entry. The detail BGP route entry is suppressed by default.

```
admin@XorPlus# set protocols bgp aggregate network6 <IPv6-Network> brief-mode <true/false>
```

Note: This command is used to create AS_SET when the BGP route entry has a different AS_PATH before aggregation.

The following examples demonstrate different methods of aggregation that are seen on the internet. The way aggregates are formed and advertised, and whether they carry more-specific routes with them will influence traffic patterns and sizes of BGP routing tables. Remember that aggregation applies to routes that exist in the BGP routing table. An aggregate can be sent if at least one more specific route of that aggregate exists in the BGP table.

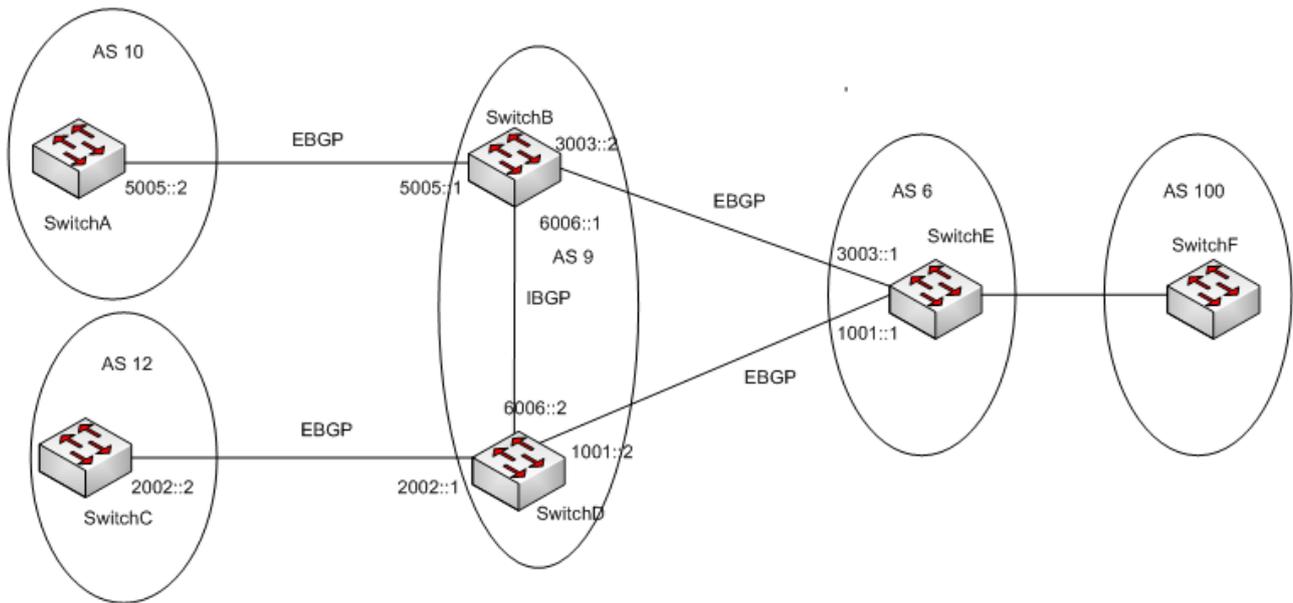
Route aggregation at border or core routers can also reduce the potential unpleasant side effects associated with IGP injection into BGP. With aggregation, multiple route entries are injected into BGP as a summary aggregate. Single route instability in any single element of the aggregate does not affect the stability of the aggregate itself.

IPv6 BGP only supports manual aggregation. Auto-summary is only for IPv4 BGP.

Aggregate Only, Suppressing the More-Specific By Default

Note: The specific route entry will be suppressed after configuring route aggregation by default.

Figure 1-22



Step 1: Configure BGP as topo showed

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 6006::2 local-ip 6006::1
admin@XorPlus# set protocols bgp peer 6006::2 as 9
admin@XorPlus# set protocols bgp peer 6006::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::1 local-ip 3003::2
admin@XorPlus# set protocols bgp peer 3003::1 as 6
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip 5005::1
admin@XorPlus# set protocols bgp peer 5005::2 as 10
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 12
admin@XorPlus# set protocols bgp peer 2002::1 local-ip 2002::2
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as 9
```

```

admin@XorPlus# set protocols bgp peer 1001::1 local-ip 1001::2
admin@XorPlus# set protocols bgp peer 1001::1 as 6
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip 2002::1
admin@XorPlus# set protocols bgp peer 2002::2 as 12
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::1 local-ip 6006::2
admin@XorPlus# set protocols bgp peer 6006::1 as 9
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true

```

Switch E:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as 6
admin@XorPlus# set protocols bgp peer 1001::2 local-ip 1001::1
admin@XorPlus# set protocols bgp peer 1001::2 as 9
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as 9
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip 4004::1
admin@XorPlus# set protocols bgp peer 4004::2 as 100
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch F:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip 4004::2
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: Switch A propagate BGP route entry 9999::/64 to Switch B. Switch C propagate BGP route entry 9999:0:0:1::/64 to Switch D.

Note: Check BGP route table on Switch E. There should be two BGP route entries from a different AS-path.

Switch E:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nextthop	Peer	AS Path
*> 9999:0:0:1::/64	3003::2	9.9.9.9	9 12 i
*> 9999::/64	1001::2	26.26.26.26	9 10 i
* 9999:0:0:1::/64	1001::2	26.26.26.26	9 12 i

```

admin@XorPlus#

```

Switch F:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nextthop	Peer	AS Path
--------	----------	------	---------

```
*> 9999:0:0:1::/64          4004::1          6.6.6.6          6 9 12 i
*> 9999::/64              4004::1          6.6.6.6          6 9 10 i
admin@XorPlus#
```

Step 3: Configure route aggregation on Switch E

Note: The specific route entry was suppressed by default. User can cancel suppress specific routes by command **set protocols bgp aggregate network6 <IPV6-Network> suppress-detail false**

Switch E:

```
admin@XorPlus# set protocols bgp aggregate network6 9999::/40
```

Check BGP route table on Switch F:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/40            4004::1                6.6.6.6             6 i
admin@XorPlus#
admin@XorPlus# run show bgp routes ipv6 detail
9999::/40
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: 6
  Nexthop: 4004::1
  Multiple Exit Discriminator: 0
  Local Preference: 100
  Aggregator: 6.6.6.6
```

Note: We can see the aggregated IPv6 network 9999::/40 on Switch F. Aggregator 6.6.6.6 created this aggregated network. The AS-path is the AS number that created this aggregated network.

Step 4: Disable suppress-detail

Note: In some cases, more-specific routes, in addition to the aggregate, need to be passed (leaked) to a neighboring AS. This is usually done in AS's multihomed to a single provider. An AS (the provider) that gets the more-specific routes would be able to make a better decision about which way to reach the route.

Switch E:

```
admin@XorPlus# set protocols bgp aggregate network6 9999::/40 suppress-detail false
```

Check the BGP route table on Switch F:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/40            4004::1                6.6.6.6             6 i
*> 9999:0:0:1::/64      4004::1                6.6.6.6             6 9 12 i
*> 9999::/64            4004::1                6.6.6.6             6 9 10 i
admin@XorPlus#
```

```

admin@XorPlus# run show bgp routes ipv6 detail
9999::/40
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 6
    Nexthop: 4004::1
    Multiple Exit Discriminator: 0
    Local Preference: 100
    Aggregator: 6.6.6.6
9999::/64
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 6 9 10
    Nexthop: 4004::1
    Multiple Exit Discriminator: 0
    Local Preference: 100
9999:0:0:1::/64
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 6 9 12
    Nexthop: 4004::1
    Multiple Exit Discriminator: 0
    Local Preference: 100

```

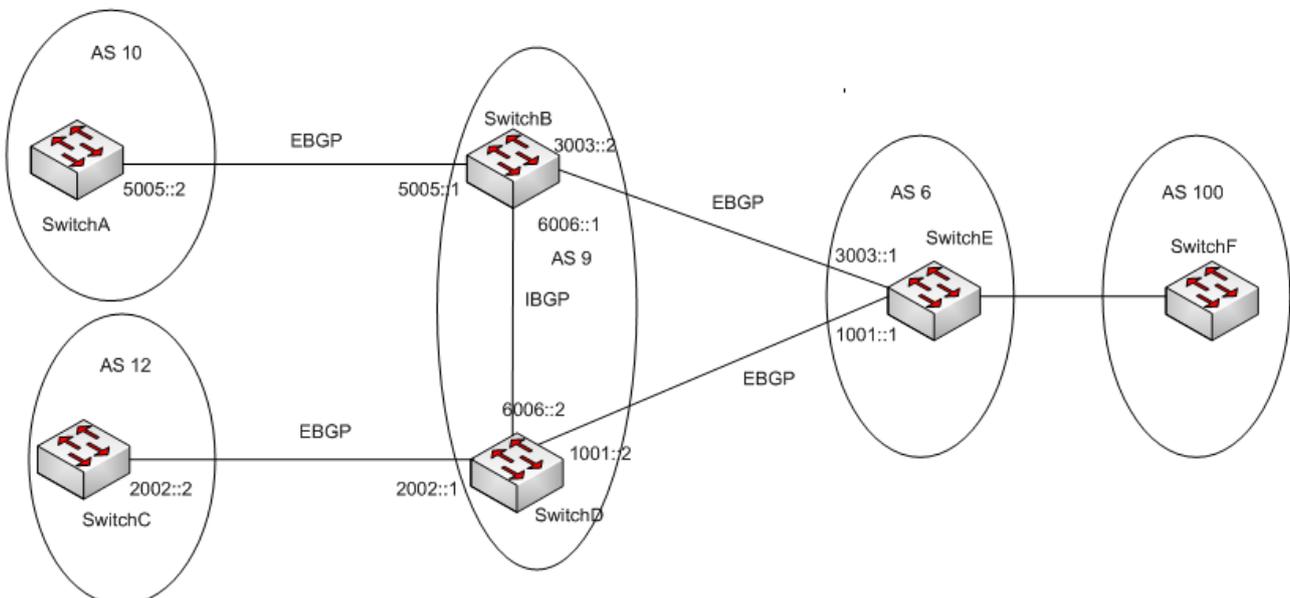
Note: The specific route entry did not suppress after disable suppress-detail.

Loss of Information Inside Aggregates

Aggregation causes loss of granularity. The detailed information that exists in the specific prefixes will be lost when summarized in the form of aggregates. The purpose of an AS_SET is to attempt to preserve the attributes carried in the specific routes in a mathematical SET that gives a better idea of the elements of the aggregate.

Note: This example demonstrates how to create as_set when the as_path is different before the BGP route entry was aggregated. It will not create as_set by default.

Figure 1-23



Step 1: Configure BGP as topo showed

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 6006::2 local-ip 6006::1
admin@XorPlus# set protocols bgp peer 6006::2 as 9
admin@XorPlus# set protocols bgp peer 6006::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::1 local-ip 3003::2
admin@XorPlus# set protocols bgp peer 3003::1 as 6
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip 5005::1
admin@XorPlus# set protocols bgp peer 5005::2 as 10
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 12
admin@XorPlus# set protocols bgp peer 2002::1 local-ip 2002::2
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as 9
admin@XorPlus# set protocols bgp peer 1001::1 local-ip 1001::2
admin@XorPlus# set protocols bgp peer 1001::1 as 6
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip 2002::1
admin@XorPlus# set protocols bgp peer 2002::2 as 12
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::1 local-ip 6006::2
admin@XorPlus# set protocols bgp peer 6006::1 as 9
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true
```

Switch E:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as 6
admin@XorPlus# set protocols bgp peer 1001::2 local-ip 1001::1
admin@XorPlus# set protocols bgp peer 1001::2 as 9
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as 9
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip 4004::1
```

```
admin@XorPlus# set protocols bgp peer 4004::2 as 100
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true
```

Switch F:

```
admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip 4004::2
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true
```

Step 2: Switch A propagate BGP route entry 9999::/64 to Switch B. Switch C propagate BGP route entry 9999:0:0:1::/64 to Switch D.

Note: Check BGP route table on Switch E. There should be two BGP route entries from different as-path.

Switch E:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix                Nexthop                Peer                AS Path
-----
*> 9999:0:0:1::/64    3003::2                9.9.9.9             9 12 i
*> 9999::/64          1001::2                26.26.26.26        9 10 i
* 9999:0:0:1::/64    1001::2                26.26.26.26        9 12 i
admin@XorPlus#
```

Switch F:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix                Nexthop                Peer                AS Path
-----
*> 9999:0:0:1::/64    4004::1                6.6.6.6             6 9 12 i
*> 9999::/64          4004::1                6.6.6.6             6 9 10 i
admin@XorPlus#
```

Step 3: Configure Route aggregation on Switch E

Note: The BGP speaker will not create AS_SET by default. User can enable the BGP speaker to create AS_SET via command **set protocols bgp aggregate network6 <IPV6-Network> brief-mode false**

Switch E:

```
admin@XorPlus# set protocols bgp aggregate network6 9999::/40
```

Check BGP route table on Switch F:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix                Nexthop                Peer                AS Path
-----
*> 9999::/40          4004::1                6.6.6.6             6 i
admin@XorPlus#
```

```
admin@XorPlus# run show bgp routes ipv6 detail
9999::/40
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 6
    Nexthop: 4004::1
    Multiple Exit Discriminator: 0
    Local Preference: 100
    Aggregator: 6.6.6.6
```

Note: The aggregated IPv6 network 9999::/40 on Switch F. Aggregator 6.6.6.6 created this aggregated network. The as-path is the AS number that created this aggregated network.

Step 4: Create AS_SET

Note: It will not Create AS_SET by default.

Switch E:

```
admin@XorPlus# set protocols bgp aggregate network6 9999::/40 brief-mode false
```

Check the BGP route table on Switch F:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::/40                            4004::1                               6.6.6.6                           6 {9 10} i
admin@XorPlus#
admin@XorPlus# run show bgp routes ipv6 detail
9999::/40
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 6 {9 10}
    Nexthop: 4004::1
    Multiple Exit Discriminator: 0
    Local Preference: 100
    Aggregator: 6.6.6.6
```

Command Reference

```
admin@XorPlus# set protocols bgp synchronization <true/false>
```

Note: This command is to enable or disable synchronization, which is disabled by default.

By definition, the default behavior of BGP requires that it must be synchronized with the IGP before BGP may advertise transit routes to external AS's. It is important that your AS be consistent about the routes it advertises, to avoid unnecessarily black-holing traffic. For example, if an IBGP speaker were to advertise a route to an external peer before all routers within your AS had learned about the route through the IGP, your AS could receive traffic to destinations for which some of the routers might not yet have the information to reach.

Whenever a router receives an update about a destination from an IBGP peer, the router tries to verify internal reachability for that destination before advertising it to other EBGP peers. The router does this by checking the destination prefix first to see if a route to the next-hop router exists and second to see if a destination prefix in the IGP exists. This router check indicates whether non-BGP routers can deliver traffic to that destination. Assuming that the IGP recognizes that destination, the router announces it to other EBGP peers. Otherwise, the router treats the destination prefix as not being synchronized with the IGP and does not advertise it.

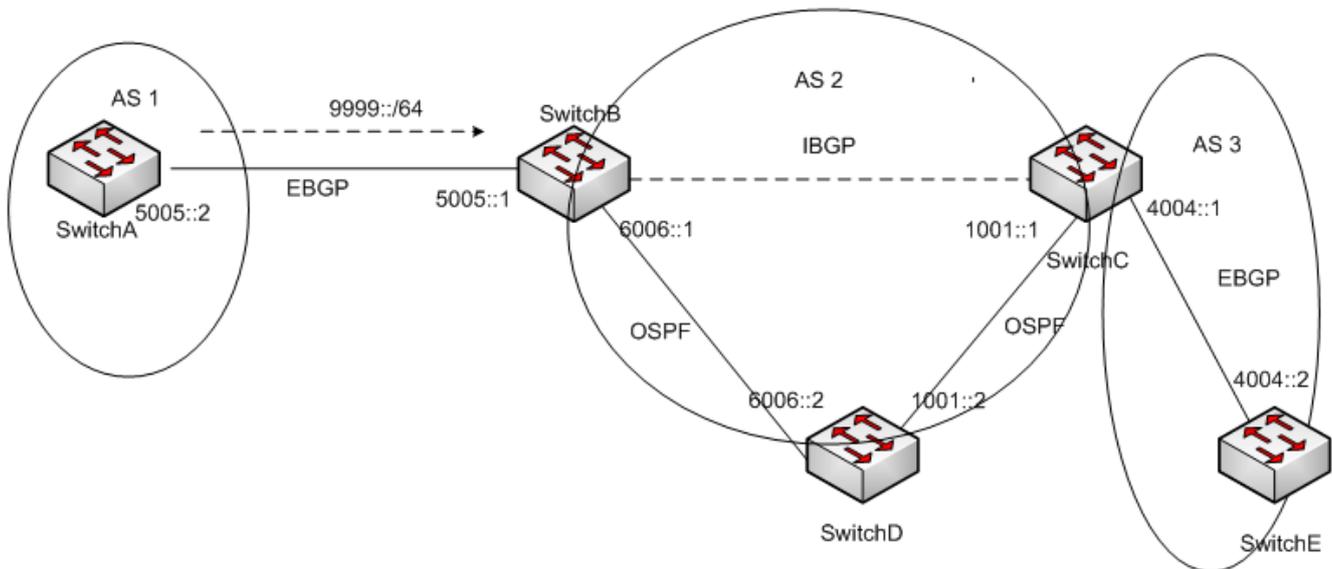
The BGP rule states that a BGP router should not advertise to external neighbors destinations learned from IBGP neighbors, unless those destinations are also known via an IGP. This is known as synchronization. If a router knows about these destinations via an IGP, it assumes that the route has already been propagated inside the AS, and internal reachability is ensured.

The consequence of injecting BGP routes inside an IGP is costly. Redistributing routes from BGP into the IGP will result in major overhead on the internal routers, primarily from an IGP scalability perspective, because (as discussed earlier) IGPs are not designed to handle that many routes. Besides, carrying all external routes inside an AS is not necessary. Routing can easily be accomplished by having internal non-BGP routers default to one of the BGP routers. Of course, this will result in suboptimal routing, because there is no guarantee that the shortest path for each route will be used, but this cost is minimal compared to maintaining thousands of routes inside the AS. Of course, managing default routes in a situation such as this can be extremely complex and may very well result in routing loops.

Most BGP implementations, however, offer a software knob that lets the network operator disable synchronization. As you might suspect, configuring **set protocols bgp synchronization false** will tell BGP to override the synchronization requirement and allow it to advertise routes learned via IBGP, irrespective of the existence of an IGP route. In practice, most situations allow synchronization to be safely turned off on border routers, assuming that all transit routers in the AS are running fully meshed IBGP. In this situation, internal reachability is guaranteed because a route that is learned via EBGP on any border router will automatically be passed on via BGP to all transit routers.

That said, by far the most common configuration in Internet-connected networks is to disable BGP synchronization, and rely on a full mesh of IBGP routers. The thought of injecting tens of thousands of routes into an IGP is quite frightening.

Figure-1 1-24



Step 1: Switch B Switch C Switch D Establish ospfv3

Switch B:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 9.9.9.9
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan600 vif vlan600 address
6006::1
```

Switch C:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 6.6.6.6
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan100 vif vlan100 address
1001::1
```

Switch D:

```
admin@XorPlus# set protocols ospf6 instance-id 1
admin@XorPlus# set protocols ospf6 router-id 26.26.26.26
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan100 vif vlan100 address
1001::2
admin@XorPlus# set protocols ospf6 area 0.0.0.0 interface vlan600 vif vlan600 address
6006::2
```

Step 2: Switch A Switch B establish EBGP, Switch B Switch C establish IGP, Switch C Switch E establish EBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 1
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2
admin@XorPlus# set protocols bgp peer 5005::1 as 2
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "2"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "6006::1"
admin@XorPlus# set protocols bgp peer 1001::1 as "2"
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "1"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "2"
admin@XorPlus# set protocols bgp peer 6006::1 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 6006::1 as "2"
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "3"
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch E:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 3
admin@XorPlus# set protocols bgp peer 4004::1 local-ip 4004::2
admin@XorPlus# set protocols bgp peer 4004::1 as 2
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 3: Enable Synchronization on Switch C

Note: If the route entry couldn't be reachable via IGP, it will not install into BGP route table.

Switch C:

```

admin@XorPlus# set protocols bgp synchronization true

```

Check Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	5005::2	33.33.33.33	1 i

```

admin@XorPlus#

```

Check Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
--------	---------	------	---------

```

admin@XorPlus#

```

Check the ospfv3 route table:

```
admin@XorPlus# run show route table ipv6 unicast ospf
6006::/64      [ospf(110)/2]
                > to fe80::ca0a:a9ff:304:4928 via vlan100/vlan100
```

Note: We can see that, the BGP route table is null, as the BGP route entry 9999::/64 didn't in IGP route table, so it will not install into BGP route table.

Step 4: Configure a policy export the BGP route entry 9999::/64 into ospfv3 on Switch B

Switch B:

```
admin@XorPlus# set policy policy-statement bgpintoospf term 1 from protocol "bgp"
admin@XorPlus# set policy policy-statement bgpintoospf term 1 then accept
admin@XorPlus# set protocols ospf6 export bgpintoospf
```

Check the ospfv3 route table on Switch C

Switch C ospfv3 route table:

```
admin@XorPlus# run show route table ipv6 unicast ospf
6006::/64      [ospf(110)/2]
                > to fe80::ca0a:a9ff:304:4928 via vlan100/vlan100
admin@XorPlus# run show route table ipv6 unicast ospf
6006::/64      [ospf(110)/2]
                > to fe80::ca0a:a9ff:304:4928 via vlan100/vlan100
9999::/64      [ospf(110)/3]
                > to fe80::ca0a:a9ff:304:4928 via vlan100/vlan100
```

Step 5: Check the BGP route table on Switch C

Note: The BGP route entry 9999::/64 should have installed into BGP route table as the route entry 9999::/64 has been in OSPFV3 route table.

Switch C BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
* 9999::/64             5005::2                9.9.9.9             1 i
```

Controlling Large-Scale Autonomous System

Route reflectors—a method of managing expanding mesh requirements in large autonomous systems (AS's) by using selected routers as focal points for internal BGP sessions.

Confederations—a method of managing expanding mesh requirements in large AS's by creating sub-AS's.

- Confederations
- Route Reflectors

Command References

```
admin@XorPlus# set protocols bgp confederation disable [true/false]
```

Note: This command is to enable/disable BGP confederation.

```
admin@XorPlus# set protocols bgp confederation identifier [confederation-identifier]
```

Note: This command is to configure a BGP confederation identifier, It's used to establish an EBGP Peer, it's a 2-byte or 4-byte AS number.

```
admin@XorPlus# set protocols bgp peer [Peer-IPv6 address] confederation-member [true/false]
```

Note: This command is to specify one BGP peer as a confederation member.

```
admin@XorPlus# set protocols bgp local-as [AS-Number]
```

Note: This Command is to configure a private local-AS number, used for establishing an internal confederation EBGP.

A confederation is another way to deal with the explosion of an IBGP mesh within an AS. As with route reflection, confederations are recommended only for cases in which IBGP peering involves a large number of IBGP peering sessions per router.

BGP confederations are based on the concept that an AS can be broken into multiple sub-AS's. Inside each sub-AS, all the rules of IBGP apply. All BGP routers inside the sub-AS, for example, must be fully meshed. Because each sub-AS has a different AS number, external BGP must run between them. Although EBGP is used between sub-ASs, routing inside the confederation behaves like IBGP routing in a single AS. In other words, the next hop, MED, and local preference information is preserved when crossing the sub-AS boundaries. To the outside world, a confederation looks like a single AS.

All the sub-ASs are shielded from the outside world and can be given any AS number. The numbers could be chosen from the private AS range (64512 to 65534, as designated in RFC 1930) in order not to use up any formal AS numbers.

As mentioned previously, inside the sub-AS, an IBGP full mesh is used. EBGP is used between the sub-AS's as well as between the confederation itself and outside AS's.

Confederations can easily detect routing loops inside the whole AS, because EBGP is run between sub-AS's. The AS path list is a loop-avoidance mechanism used to detect routing updates leaving one sub-AS and attempting to reenter the same sub-AS. A routing update that tries to reenter the sub-AS it originated from will be detected, because the sub-AS will see its own sub-AS number listed in the update's AS path.

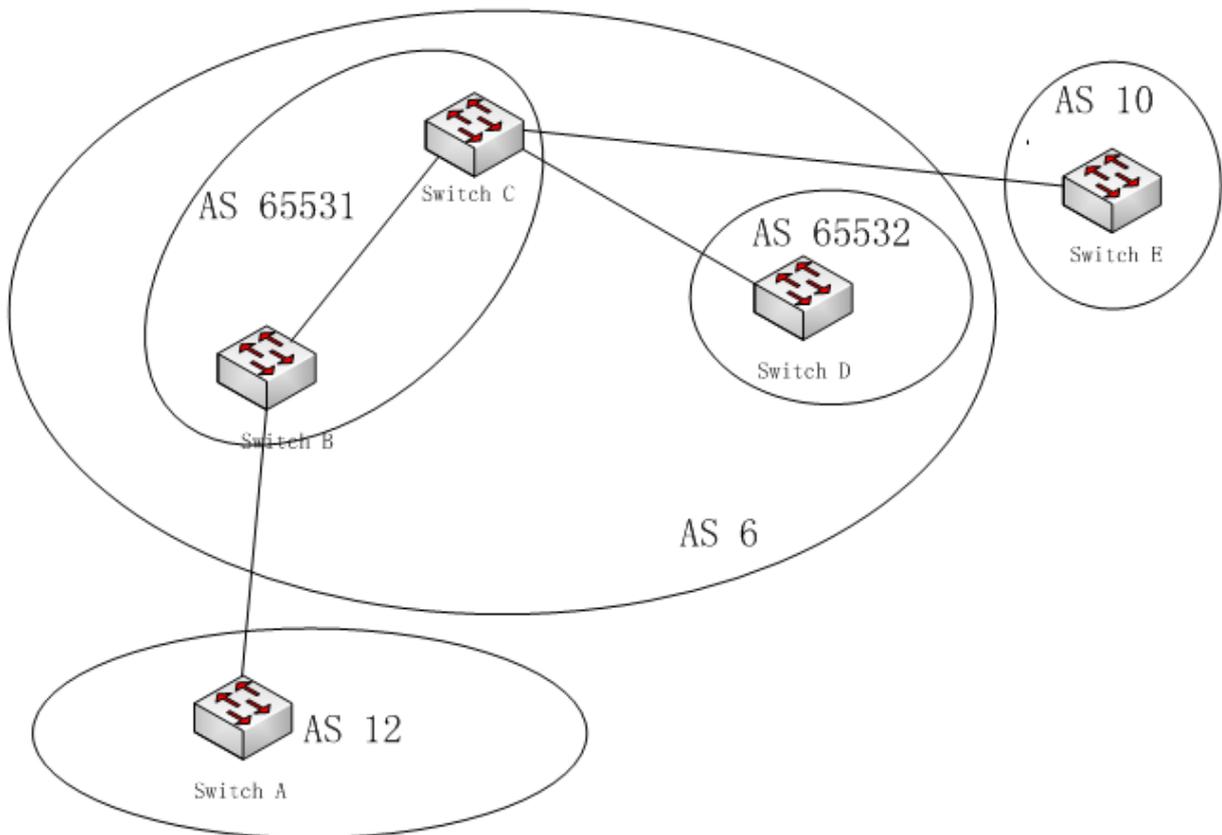


Figure 2-5

Step 1: Switch B, Switch C, and Switch D configured as a confederation member

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "65531"
admin@XorPlus# set protocols bgp confederation identifier "6"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "65531"
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 confederation-member true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "65531"
admin@XorPlus# set protocols bgp confederation identifier "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "65532"
admin@XorPlus# set protocols bgp peer 1001::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::2 client true
admin@XorPlus# set protocols bgp peer 1001::2 confederation-member true
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "65531"
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::2 confederation-member true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch D:

```

admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "65532"
admin@XorPlus# set protocols bgp confederation identifier "6"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "65531"
admin@XorPlus# set protocols bgp peer 1001::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::1 confederation-member true
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true

```

Step 2: Switch A and Switch B configure EBGP, Switch C and Switch E configure EBGP

Switch A:

```

admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip "5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 6
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true

```

Switch D:

```

admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as 100
admin@XorPlus# set protocols bgp peer 4004::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch E:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 3: Switch A distribute one BGP route entry, then check the BGP route table on Switch B, Switch C, Switch D and Switch E.

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	5005::2	33.33.33.33	10 i

```

admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 33.33.33.33

```

```

Route: Winner
Origin: IGP
AS Path: 10
Nexthop: 5005::2
Local Preference: 100
admin@XorPlus#

```

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    3003::2         9.9.9.9       10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 3003::2
  Local Preference: 100
admin@XorPlus#

```

Switch D BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    1001::1         6.6.6.6       (65531) 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: (65531) 10
  Nexthop: 1001::1
  Local Preference: 100
admin@XorPlus#

```

Note: We can see that the AS-path includes the confederation EGBP AS number 65531

Switch E BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    4004::1         6.6.6.6       6 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: 6 10
  Nexthop: 4004::1
  Local Preference: 100
admin@XorPlus#

```

Note: We can see that the confederation AS number 65531 was removed from the AS-path when it outbounds the AS to it's EBGP Peer.

Route Reflectors Command References

```
admin@XorPlus# set protocols bgp route-reflector disable [true/false]
```

Note: This command is to enable/disable route-reflectors.

```
admin@XorPlus# set protocols bgp route-reflector cluster-id [cluster-id]
```

Note: This command is to specify a cluster-id.

```
admin@XorPlus# set protocols bgp peer [bgp peer id] client [true/false]
```

Note: This command is to specify one BGP peer as RR's client.

In some Internet service provider (ISP) networks, the internal BGP mesh can become quite large (more than 100 internal BGP sessions per router), which strongly suggests that some new peering mechanism should be implemented. The *route reflector* concept is based on the idea of specifying a *concentration* router to act as a focal point for internal BGP sessions. Multiple (client) BGP routers can peer with a central server (the route reflector), and then route reflectors peer with one another. Although the BGP rule states that routes learned via one IBGP speaker can't be advertised to another IBGP speaker, route reflection allows the route reflector servers to "reflect" routes as described later, thereby relaxing the IBGP full-mesh constraints.

The route reflector is a router that performs the route reflection function. The IBGP peers of the route reflector fall under two categories—*clients* and *nonclients*. A route reflector and its clients form a *cluster*. All peers of the route reflector that are not part of the cluster are nonclients.

Nonclients (standard IBGP speakers) are still required to be fully meshed with one another and the route reflector, because they follow the normal IBGP advertisement rules, although they no longer need to peer with the clients of the route reflectors. Clients should not peer with internal speakers outside their associated cluster.

The route reflector function is implemented only on the route reflector; all clients and nonclients are normal BGP peers that have no notion of the route reflector. Route reflector clients are considered as such only because the route reflector lists them as clients.

Any route reflector that receives multiple routes for the same destination employs the usual BGP decision process to pick the overall best path. The best path would be propagated inside the AS based on the following rules of operation:

- If the route is received from a nonclient peer, reflect to clients only.
- If the route is received from a client peer, reflect to all nonclient peers and also to client peers.
- If the route is received from an EBGP peer, reflect to all client and nonclient peers.

Because route reflection is a concept that applies only internally to an AS, routers external to the AS, which would receive UPDATEs via EBGP, are considered nonclients and follow normal nonclient behavior with respect to sending and receiving UPDATEs.

The Route Reflector Preserves IBGP Attributes

The route reflector concept does not change IBGP behavior—the route reflector is not allowed to modify the attributes of the reflected IBGP routes. The NEXT_HOP attribute, for example, remains the same when an IBGP route is exchanged between RRs. This is necessary to avoid loops inside the AS.

Avoiding Loops

BGP relies on the information in the AS path to facilitate loop detection. A BGP update that attempts to reenter the AS it was originated from, will be dropped by the border router of the source AS. With the introduction of route reflectors, there is a potential for routing loops within an AS. A routing update that leaves a cluster may reenter the cluster. Loops inside the AS cannot be detected by the traditional AS path approach because routing updates do not have an originating AS path signature. Therefore, when route reflectors are deployed, BGP offers two extra measures for loop avoidance inside the AS—using an ORIGINATOR_ID and using a CLUSTER_LIST.

Using an ORIGINATOR_ID

The ORIGINATOR_ID is a 4-byte, optional, nontransitive BGP attribute (type code 9). This attribute carries the ROUTER_ID of the route's originator in the local AS and is to be added to the UPDATE message by the route reflector. If the update comes back to the originator because of poor configuration, the originator should discard it.

The CLUSTER_LIST

The CLUSTER_LIST is an optional, nontransitive BGP attribute (type code 10). Each cluster is represented with a CLUSTER_ID. A CLUSTER_LIST is a sequence of CLUSTER_IDs that contain path information regarding the list of clusters that an UPDATE has traversed. When a route reflector sends a route from its clients to nonclients outside the cluster, it appends the local CLUSTER_ID to the CLUSTER_LIST, or creates the list if one is not present. If the route reflector receives an UPDATE whose CLUSTER_LIST contains the local CLUSTER_ID value, the UPDATE message should be discarded. Thus, the CLUSTER_LIST provides loop avoidance inside an AS, whereas the AS_PATH list, discussed earlier, facilitates loop avoidance for UPDATEs traversing multiple, external ASs.

Configuration Example 1:

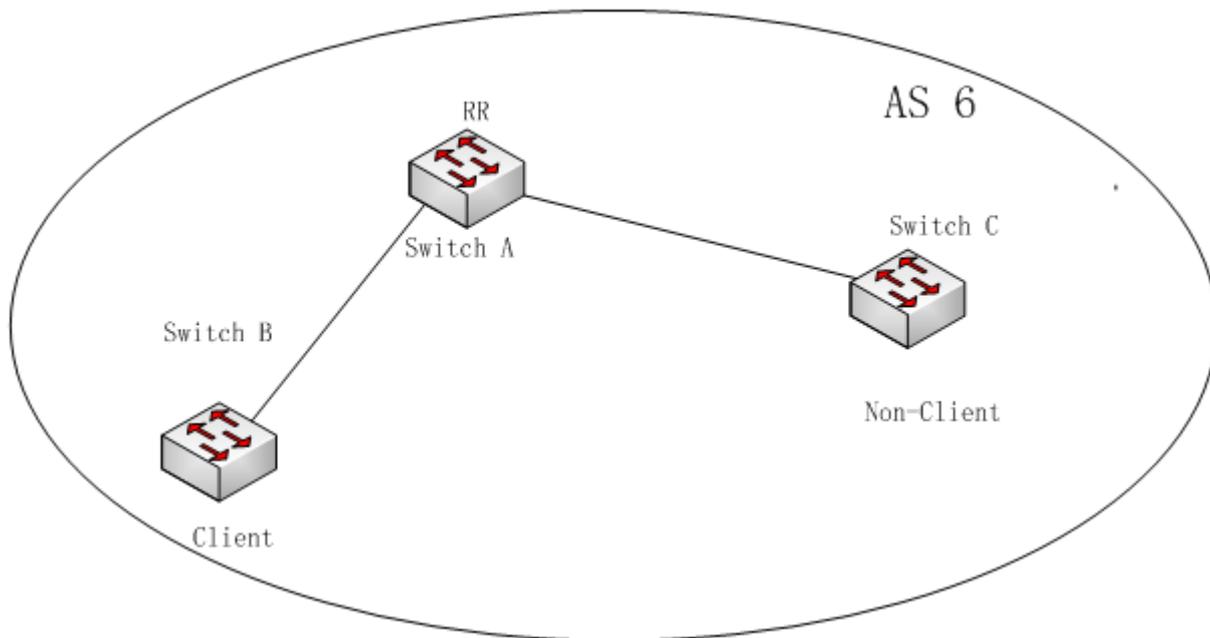


Figure 2-1

Step 1: Configure IBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "6"
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Step 2: Configure Switch A as RR and configure Cluster-ID

```
admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp route-reflector cluster-id 6.6.6.6
```

Step 3: Specify one BGP peer as RR's client

```
admin@XorPlus# set protocols bgp peer 3003::2 client true
```

Step 4: Check BGP peer status:

Switch A:

```
admin@XorPlus# run show bgp peers detail
Peer 2: local 3003::1/179 remote 3003::2/33239
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 1, Updates Sent: 6
Messages Received: 102, Messages Sent: 108
Time since last received update: 2611 seconds
Number of transitions to ESTABLISHED: 4
Time since last entering ESTABLISHED state: 2611 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Switch B:

```
admin@XorPlus# run show bgp peers detail
Peer 1: local 3003::2/33239 remote 3003::1/179
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 6, Updates Sent: 1
Messages Received: 123, Messages Sent: 118
Time since last received update: 1079 seconds
Number of transitions to ESTABLISHED: 9
Time since last entering ESTABLISHED state: 2988 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
```

Configuration Example 2:

BGP Reflector Configuration examples.

RR will propagate BGP route to all it's Client and Non-client peer:

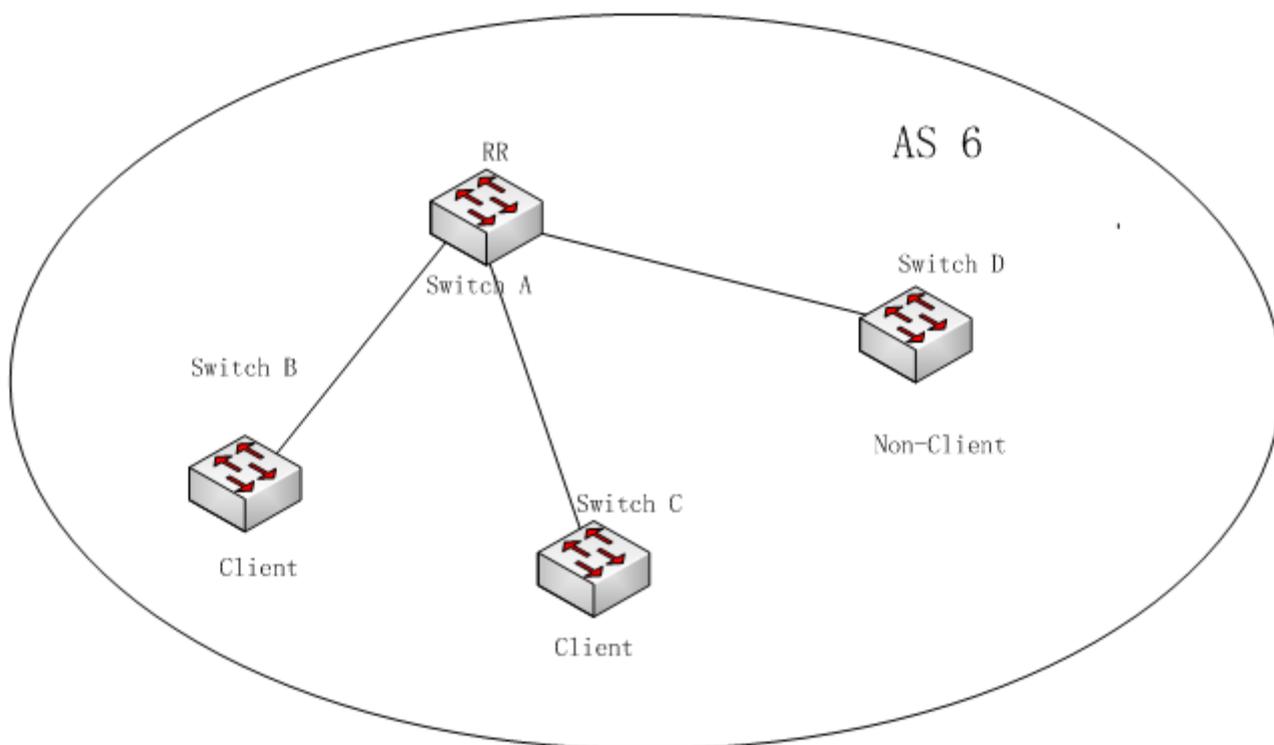


Figure 2-2

Step 1: Configure IBGP on Switch A, Switch B, Switch C, Switch D

Switch A:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "6"
admin@XorPlus# set protocols bgp peer 1001::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "6"
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "6"
admin@XorPlus# set protocols bgp peer 4004::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 10.10.10.10
admin@XorPlus# set protocols bgp local-as "6"

```

```
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as "6"
admin@XorPlus# set protocols bgp peer 4004::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
```

Step 2: Configure Switch A as RR and configure Cluster-ID

```
admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp route-reflector cluster-id 6.6.6.6
```

Step 3: Specify Switch B and Switch C as RR's Client

```
admin@XorPlus# set protocols bgp peer 3003::2 client true
admin@XorPlus# set protocols bgp peer 4004::2 client true
```

Step 4: Switch B distributes a BGP route, then RR will propagate to all it's clients (Switch C) and non-client (Switch D)

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    5005::2         33.33.33.33   10 i
Switch A bgp route table:
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    3003::2         9.9.9.9       10 i
```

Check the BGP route table of Switch C, we can see that the Client got a BGP route entry from RR.

Switch C BGP route table:

```
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 4004::1
  Multiple Exit Discriminator: 1000
  Local Preference: 100
  Originator ID: 9.9.9.9
  Cluster List: 6.6.6.6
```

Check the BGP route table of Switch D, we can see that Non-Client (Switch D) got a BGP route entry from RR

Switch D BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    1001::1         6.6.6.6      10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 1001::1
  Multiple Exit Discriminator: 1000
  Local Preference: 100
  Originator ID: 9.9.9.9
  Cluster List: 6.6.6.6
```

Configuration Example 3:

2. A route from a non-client peer is advertised to all clients. It means the route entry from a non-client peer will not be advertised to RR's non-client peer:

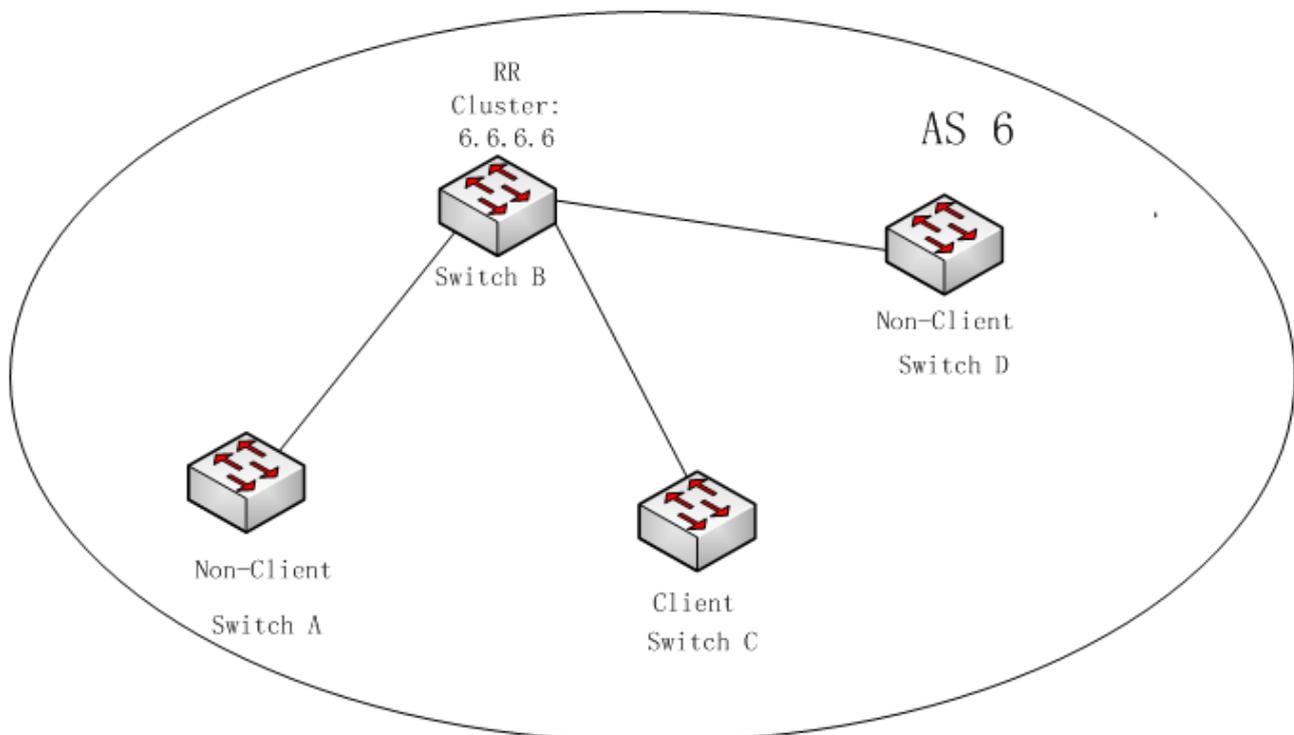


Figure 2-3

Step 1: Configure IBGP on SwitchA SwitchB SwitchC SwitchD

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "6"
```

```

admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "6"
admin@XorPlus# set protocols bgp peer 1001::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "6"
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "6"
admin@XorPlus# set protocols bgp peer 4004::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true

```

Switch D:

```

admin@XorPlus# set protocols bgp bgp-id 10.10.10.10
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as "6"
admin@XorPlus# set protocols bgp peer 4004::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: Enable route-reflector on Switch B and specify Switch C as it's client

```

admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp route-reflector cluster-id 6.6.6.6
admin@XorPlus# set protocols bgp peer 1001::2 client true

```

Step 3: Switch A distributes two BGP route entries, then check the BGP route table on Switch B, Switch C, Switch D, we can see that Switch B will not advertise these BGP route entries to it's non-client peer (Switch D)

Switch B BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	3003::2	9.9.9.9	10 i

```
*> 9999:0:0:1::/64          3003::2          9.9.9.9          10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 10
    Nexthop: 3003::2
    Local Preference: 100
9999:0:0:1::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 10
    Nexthop: 3003::2
    Local Preference: 100
admin@XorPlus#
```

Switch C BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
*> 9999::/64    1001::1         6.6.6.6      10 i
*> 9999:0:0:1::/64 1001::1         6.6.6.6      10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 10
    Nexthop: 1001::1
    Local Preference: 100
    Originator ID: 9.9.9.9
    Cluster List: 6.6.6.6
9999:0:0:1::/64
    From peer: 6.6.6.6
    Route: Winner
    Origin: IGP
    AS Path: 10
    Nexthop: 1001::1
    Local Preference: 100
    Originator ID: 9.9.9.9
    Cluster List: 6.6.6.6
admin@XorPlus#
```

Switch D BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
Prefix          Nexthop          Peer          AS Path
-----
```

Note: The BGP route table on Switch D should be NULL, as the BGP route entry from non-client will not advertise to its non-client peer.

Configuration Example 4:

3. Multistage bgp reflector:

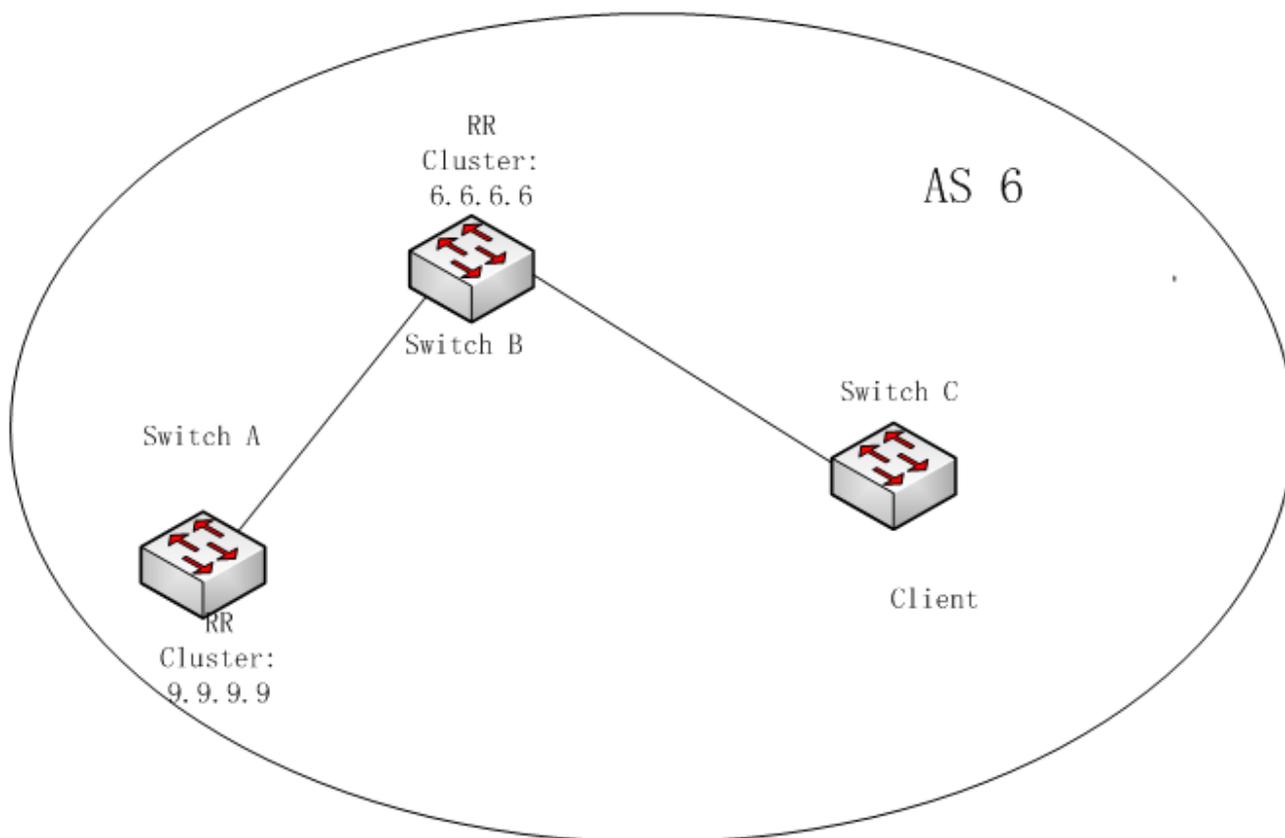


Figure 2-4

Step 1: Configure IBGP on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "6"
admin@XorPlus# set protocols bgp peer 1001::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "6"
admin@XorPlus# set protocols bgp peer 3003::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
```

```
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
```

Step 2: Configure Switch A as first-level RR ,Switc hB is Switch A's Client

Switch A:

```
admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp route-reflector cluster-id 9.9.9.9
admin@XorPlus# set protocols bgp peer 3003::1 client true
```

Step 3:Configure Switch B as the second-level RR,Switch C is Switch B's Client

Switch B:

```
admin@XorPlus# set protocols bgp route-reflector disable false
admin@XorPlus# set protocols bgp route-reflector cluster-id 6.6.6.6
admin@XorPlus# set protocols bgp peer 1001::2 client true
```

Step 4:Switch A got a BGP route entry from an EBGP peer, then check the BGP route table on Switch A, Switch B, and Switch C,

Switch A BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            5005::2                33.33.33.33         10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 33.33.33.33
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 5005::2
  Multiple Exit Discriminator: 1000
  Local Preference: 100
admin@XorPlus#
```

Check the BGP route table of Switch B,we can see the router-id of Switch A

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2                9.9.9.9             10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 3003::2
  Multiple Exit Discriminator: 1000
```

```

Local Preference: 100
Originator ID: 33.33.33.33
Cluster List: 9.9.9.9
admin@XorPlus#

```

Check the BGP route table of Switch C, we can see that the router-id of Switch B was add into cluster list

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64           1001::1                6.6.6.6            10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 6.6.6.6
  Route: Winner
  Origin: IGP
  AS Path: 10
  Nexthop: 1001::1
  Multiple Exit Discriminator: 1000
  Local Preference: 100
  Originator ID: 33.33.33.33
  Cluster List: 6.6.6.6, 9.9.9.9
admin@XorPlus#

```

BGP Load Balance Command References

```

admin@XorPlus# set protocols bgp multipath disable [true/false]

```

Note: This command is to enable/disable Load Balance when it have multiple connections to the same destination and all the BGP route entry have the same AS-path values.

```

admin@XorPlus# set protocols bgp multipath path-relax [true/false]

```

Note: This command is to enable/disable path-relax, when it have multiple connections to the same destination and they have the different AS-path, if you enable path-relax,they also could form Load Balance.

Under normal conditions, when a BGP speaker receives identical paths for a prefix from an adjacent AS, only one path will be selected as the best path (normally the one with the lowest BGP ROUTER_ID value) and will be installed in the routing table. If BGP multipath is enabled, multiple paths can be installed in the IP routing table. By default, just all the BGP routes have the same AS-paths, they can form Load Balance, but if you enable path-relax, the BGP route have the different AS-paths, they also can form Load Balance.

Configure Example 1: Redundancy

When a BGP speaker receives identical paths for a prefix from an adjacent AS, only one path will be selected as the best path and will be installed in the routing table. When the primary failed, the backup one will be selected as the best path and will be installed in the routing table.

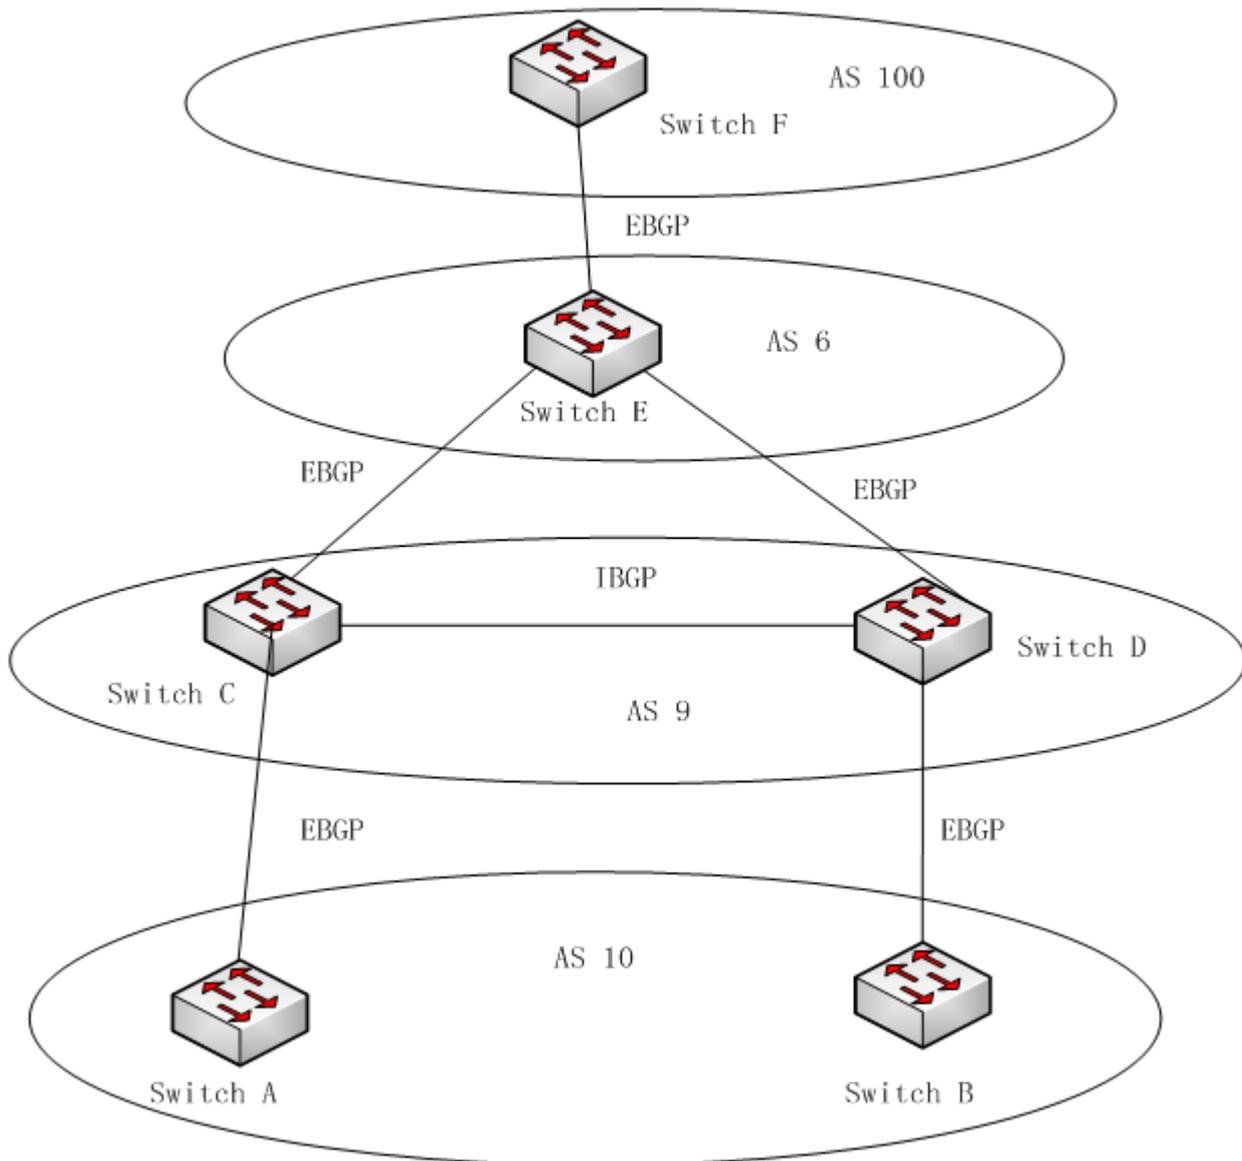


Figure 3-1

Step 1: Switch A, Switch C configure EBGP, Switch B, Switch D configure EBGP, Switch C. Switch D configure IBGP, Switch C, Switch E configure EBGP, Switch D, Switch E configure EBGP, Switch E, Switch F configure EBGP, Switch A, Switch B in AS 10, Switch C, Switch D in AS 9, Switch E in as 6, Switch F in AS 100

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
```

```
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 2002::1 local-ip 2002::2"
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::2 local-ip "6006::1"
admin@XorPlus# set protocols bgp peer 6006::2 as "9"
admin@XorPlus# set protocols bgp peer 6006::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::2 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip "2002::1"
admin@XorPlus# set protocols bgp peer 2002::2 as "10"
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::1 local-ip "6006::2"
admin@XorPlus# set protocols bgp peer 6006::1 as "9"
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true
```

Switch E:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "9"
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "100"
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true
```

Switch F:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: Switch A and Switch B distribute BGP route entry 9999::/64, then check the BGP route table on Switch C, Switch D, Switch E

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	5005::2	33.33.33.33	10 i
* 9999::/64	6006::2	26.26.26.26	10 i

Switch D BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	6006::1	9.9.9.9	10 i
*> 9999::/64	2002::2	44.44.44.44	10 i

Switch E BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	3003::2	9.9.9.9	9 10 i
* 9999::/64	1001::2	26.26.26.26	9 10 i

Step 3: The primary link down, then the BGP route entry from the backup link will be the best route, then check the BGP route table on Switch E

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	1001::2	26.26.26.26	9 10 i

Configuration Example 2 : Load Balance (BGP route have the same AS-path)

When a BGP speaker receives identical paths for a prefix from an adjacent AS, if BGP multipath is enabled, multiple paths can be installed in the IP routing table.

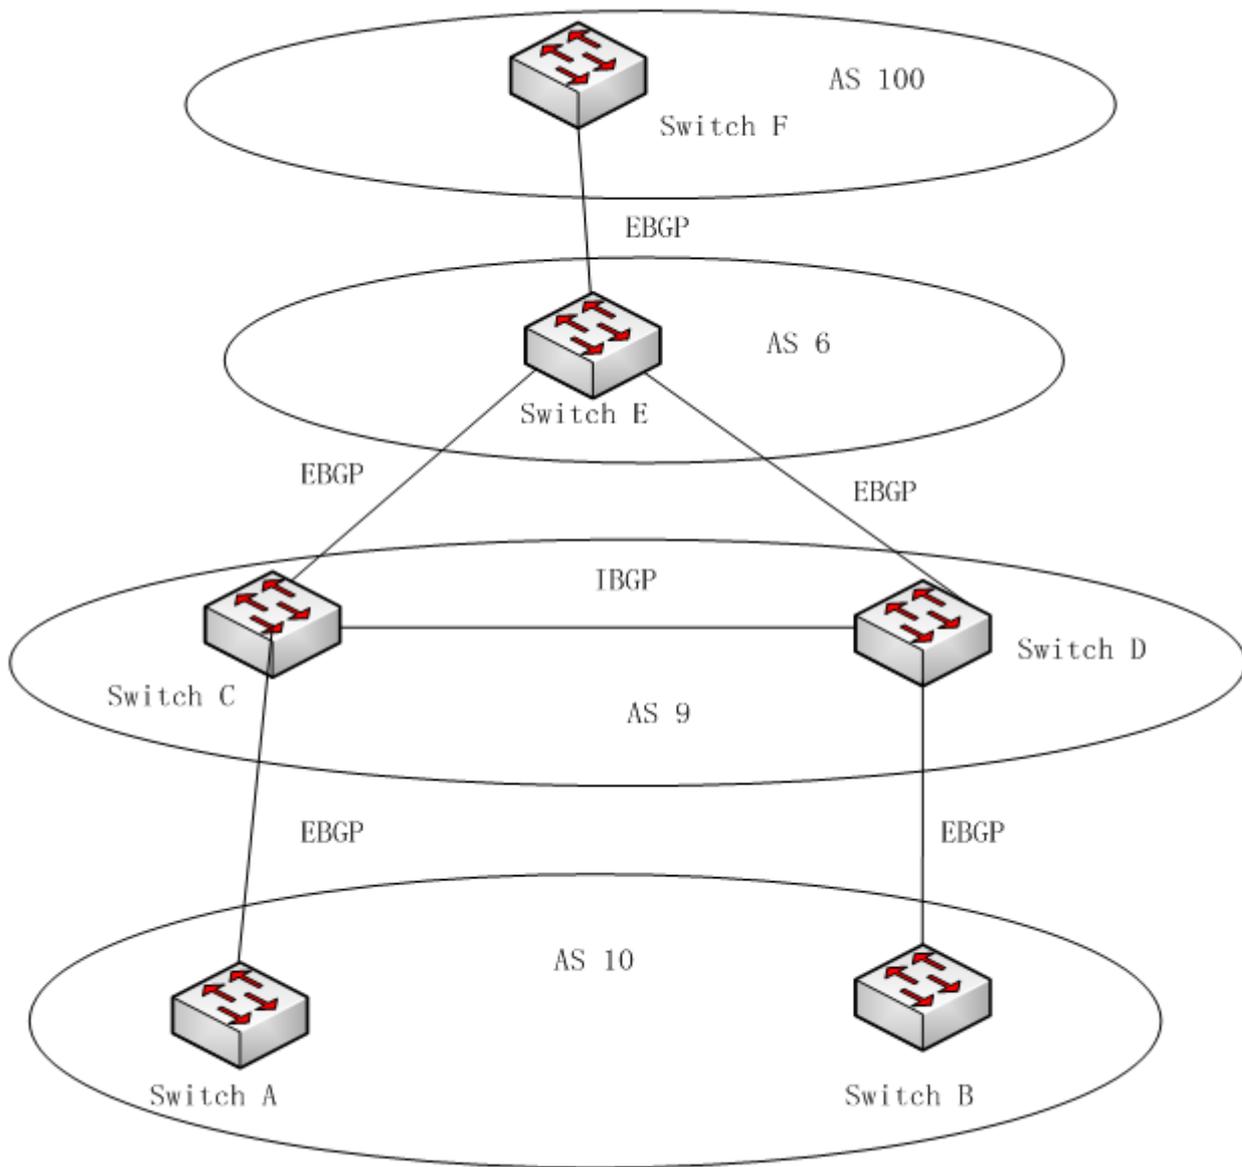


Figure 3-2

Step 1: Switch A Switch C configure EBGP, Switch B Switch D configure EBGP, Switch C Switch D configure IBGP, Switch C Switch E configure EBGP, Switch D Switch E configure EBGP, Switch E Switch F configure EBGP, Switch A Switch B in AS 10, Switch C Switch D in AS 9, Switch E in AS 6, Switch F in AS 100

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 10
admin@XorPlus# set protocols bgp peer 2002::1 local-ip 2002::2"
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::2 local-ip "6006::1"
admin@XorPlus# set protocols bgp peer 6006::2 as "9"
admin@XorPlus# set protocols bgp peer 6006::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::2 ipv6-unicast true

```

Switch D:

```

admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip "2002::1"
admin@XorPlus# set protocols bgp peer 2002::2 as "10"
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::1 local-ip "6006::2"
admin@XorPlus# set protocols bgp peer 6006::1 as "9"
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true

```

Switch E:

```

admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "9"
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "100"
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true

```

Switch F:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: SwitchA and SwitchB distribute BGP route entry 9999::/64, then check the BGP route table on Switch C, Switch D, SwitchE.

Switch C BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::/64                            5005::2                               33.33.33.33                       10 i
* 9999::/64                             6006::2                               26.26.26.26                       10 i
```

Switch D BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
* 9999::/64                             6006::1                               9.9.9.9                           10 i
*> 9999::/64                             2002::2                               44.44.44.44                       10 i
```

Switch E BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::/64                            3003::2                               9.9.9.9                           9 10 i
* 9999::/64                             1001::2                               26.26.26.26                       9 10 i
```

Step 3: Enable Multipath on Switch E, then check BGP route table, it will form Load Balance, as the two BGP route entry have the same AS-path

```
admin@XorPlus#set protocols bgp multipath disable false
```

Switch E BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::/64                            1001::2                               26.26.26.26                       9 10 i
*> 9999::/64                             3003::2                               9.9.9.9                           9 10 i
admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                               Port
-----                               -
2001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
4004::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
1001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
3003::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF connected
9999::                                     ffff:ffff:ffff:ffff::
```

```

C8:0A:A9:AE:0A:66   te-1/1/46
9999::              ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28   te-1/1/48

```

Configuration Example 3 : Load Balance (BGP route have different AS-path)

When a BGP speaker receives identical paths for a prefix from an adjacent AS, if BGP multipath is enabled, multiple paths can be installed in the IP routing table:

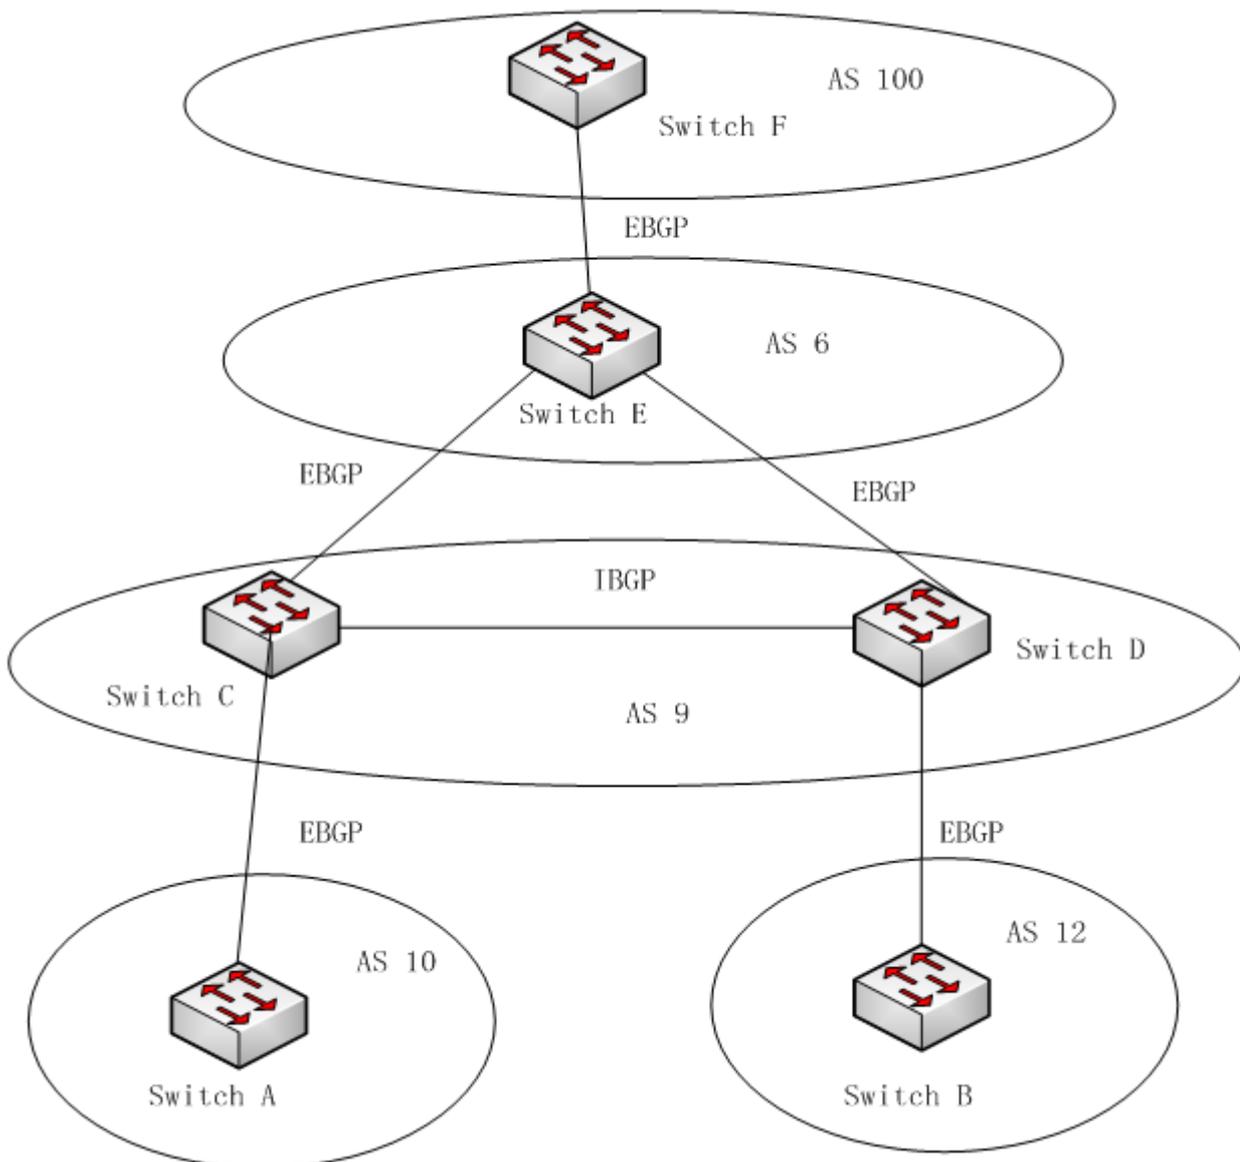


Figure 3-3

Step 1: Switch A, Switch C configure EBGP, Switch B, Switch D configure EBGP, Switch C, Switch D configure IBGP, Switch C, Switch E configure EBGP, Switch D, Switch E configure EBGP, Switch E, Switch F configure EBGP, Switch A in AS 10, Switch B in AS 12, Switch C, Switch D in AS 9, Switch E in AS 6, Switch F in AS 100

Switch A:

```

admin@XorPlus# set protocols bgp bgp-id 33.33.33.33
admin@XorPlus# set protocols bgp local-as 10

```

```
admin@XorPlus# set protocols bgp peer 5005::1 local-ip 5005::2"
admin@XorPlus# set protocols bgp peer 5005::1 as 9
admin@XorPlus# set protocols bgp peer 5005::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 44.44.44.44
admin@XorPlus# set protocols bgp local-as 12
admin@XorPlus# set protocols bgp peer 2002::1 local-ip 2002::2"
admin@XorPlus# set protocols bgp peer 2002::1 as 9
admin@XorPlus# set protocols bgp peer 2002::1 ipv6-unicast true
```

Switch C:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 5005::2 local-ip "5005::1"
admin@XorPlus# set protocols bgp peer 5005::2 as "10"
admin@XorPlus# set protocols bgp peer 5005::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::2 local-ip "6006::1"
admin@XorPlus# set protocols bgp peer 6006::2 as "9"
admin@XorPlus# set protocols bgp peer 6006::2 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::2 ipv6-unicast true
```

Switch D:

```
admin@XorPlus# set protocols bgp bgp-id 26.26.26.26
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 1001::1 local-ip "1001::2"
admin@XorPlus# set protocols bgp peer 1001::1 as "6"
admin@XorPlus# set protocols bgp peer 1001::1 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 2002::2 local-ip "2002::1"
admin@XorPlus# set protocols bgp peer 2002::2 as "12"
admin@XorPlus# set protocols bgp peer 2002::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 6006::1 local-ip "6006::2"
admin@XorPlus# set protocols bgp peer 6006::1 as "9"
admin@XorPlus# set protocols bgp peer 6006::1 next-hop-self true
admin@XorPlus# set protocols bgp peer 6006::1 ipv6-unicast true
```

Switch E:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 1001::2 local-ip "1001::1"
admin@XorPlus# set protocols bgp peer 1001::2 as "9"
admin@XorPlus# set protocols bgp peer 1001::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
admin@XorPlus# set protocols bgp peer 4004::2 local-ip "4004::1"
admin@XorPlus# set protocols bgp peer 4004::2 as "100"
admin@XorPlus# set protocols bgp peer 4004::2 ipv6-unicast true
```

Switch F:

```

admin@XorPlus# set protocols bgp bgp-id 100.100.100.100
admin@XorPlus# set protocols bgp local-as 100
admin@XorPlus# set protocols bgp peer 4004::1 local-ip "4004::2"
admin@XorPlus# set protocols bgp peer 4004::1 as 6
admin@XorPlus# set protocols bgp peer 4004::1 ipv6-unicast true

```

Step 2: Switch A, Switch B distribute BGP route entry 9999::/64, then check the BGP route table on Switch C, Switch D, Switch E

Switch C BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	6006::2	26.26.26.26	12 i
*> 9999::/64	5005::2	33.33.33.33	10 i

Switch D BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
* 9999::/64	6006::1	9.9.9.9	10 i
*> 9999::/64	2002::2	44.44.44.44	12 i

Switch E BGP route table:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	3003::2	9.9.9.9	9 10 i
* 9999::/64	1001::2	26.26.26.26	9 12 i

Step 3: Enable Multipath and enable path-relax on Switch E (as the two BGP route have different AS-path, if it want to form Load Balance, it must enable path-relax and enable multipath at the same time), then check the BGP route table on Switch E, it will form Load Balance.

```

admin@XorPlus# set protocols bgp multipath disable false
admin@XorPlus# set protocols bgp multipath path-relax true
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 9999::/64	1001::2	26.26.26.26	9 12 i
*> 9999::/64	3003::2	9.9.9.9	9 10 i

```

admin@XorPlus# run show route forward-route ipv6 all
Destination                               NetMask
NextHopMac                                Port
-----
2001::                                     ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF                         connected

```

```

4004::                                ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF  connected
1001::                                ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF  connected
3003::                                ffff:ffff:ffff:ffff::
04:7D:7B:62:93:FF  connected
9999::                                ffff:ffff:ffff:ffff::
C8:0A:A9:AE:0A:66  te-1/1/46
9999::                                ffff:ffff:ffff:ffff::
C8:0A:A9:04:49:28  te-1/1/48

```

Step 4: Check the BGP route table on Switch F, the BGP route entry was gotten from Switch E, and it's the first one that arrived at Switch E

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                               Nexthop                               Peer                               AS Path
  -----                               -
*> 9999::/64                           4004::1                               6.6.6.6                           6 9 12 i

```

Route Dampening Command Reference

```
admin@XorPlus# set protocols bgp damping disable [true/false]
```

Note: This command is to Enable/Disable route dampening.

```
admin@XorPlus# set protocols bgp damping half-life [1-45 minutes]
```

Note: This command is to set the A configurable numeric value that describes the amount of time that must elapse to reduce the penalty by one-half, the default value is 15 minutes.

```
admin@XorPlus# set protocols bgp damping max-suppress [1-720 minutes]
```

Note: This command is to set the max-suppress timer, it's default value is 60 minutes.

```
admin@XorPlus# set protocols bgp damping reuse [1-20000 seconds]
```

Note: This command is to set A configurable numeric value that is compared with the penalty. If the penalty is less than the reuse limit, a suppressed route that is up will no longer be suppressed, the default value is 750.

```
admin@XorPlus# set protocols bgp damping suppress [1-20000 seconds]
```

Note: This command is to set a numeric value that is compared with the penalty. If the penalty is greater than the suppress limit, the route is suppressed; the default value is 3000.

Another mechanism for controlling route instability is *route dampening*. A route that appears and disappears intermittently causes BGP UPDATE and WITHDRAWN messages to be repeatedly propagated on the Internet. The tremendous amount of routing traffic generated can use up all the link's bandwidth and drive up CPU utilization of routers.

Dampening categorizes routes as well either *behaved* or *ill behaved*. A well-behaved route shows a high degree of stability during an extended period of time. On the other hand, an ill behaved route experiences a high level of instability in a short period of time. Ill-behaved routes should be penalized in a way that is proportional to the route's expected future instability. An unstable route should be suppressed (not advertised) until there is some degree of confidence that the route has become stable.

A route's recent history is used as a basis for estimating future stability. To track a route history, it is essential to track the number of times the route has flapped over a period of time. Under route dampening, each time a route flaps, it is given a penalty. Whenever the penalty reaches a predefined threshold, the route is suppressed. The route can continue to accrue penalties even after it is suppressed. The more frequently a route oscillates in a short amount of time, the faster the route is suppressed.

Similar criteria are put in place to un-suppress a route and start readvertising it. An algorithm is implemented to decay (reduce) the penalty value exponentially.

Penalty

An incremented numeric value that is assigned to a route each time it flaps.

Half-life

A configurable numeric value that describes the amount of time that must elapse to reduce the penalty by one-half.

Suppress Limit

A numeric value that is compared with the penalty. If the penalty is greater than the suppress limit, the route is suppressed.

Reuse Limit

A configurable numeric value that is compared with the penalty. If the penalty is less than the reuse limit, a suppressed route that is up will no longer be suppressed.

Max-suppress Timer

It is the maximum time the BGP route was suppressed.

The following figure illustrates the process of assessing a penalty to a route every time it flaps. The penalty is exponentially decayed according to parameters such as the half-life. The half-life parameter can be changed by the administrator to reflect the oscillation history of a route: A longer half-life might be desirable for a route that has a habit of oscillating frequently. A larger half-life value would cause the penalty to decay more slowly, which translates into a route's being suppressed longer.

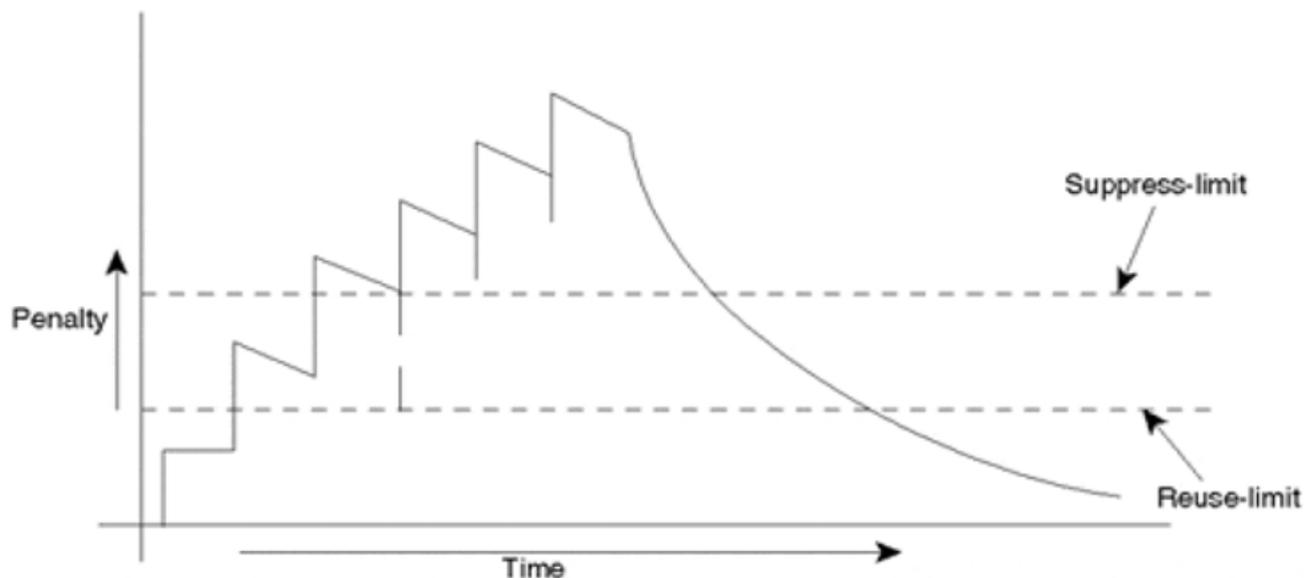


Figure 4-1 Route Dampening Penalty Assessment

Configuration Example 1:

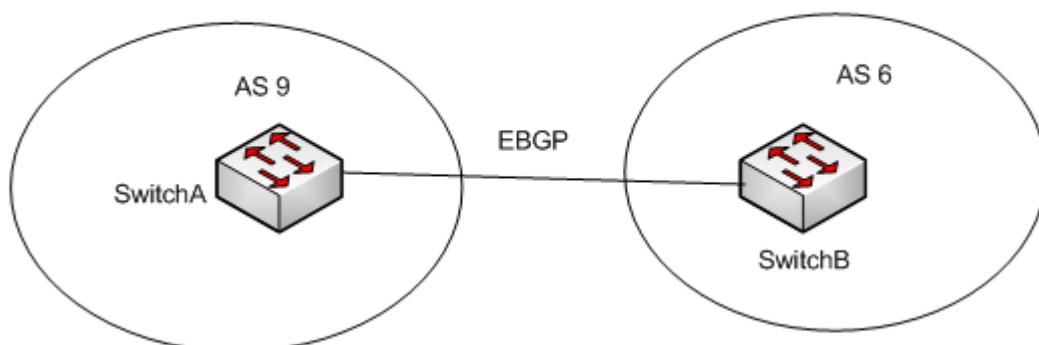


Figure 4-2

Step 1: Switch A, Switch B Establishing EBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Step 2: Switch A propagate a BGP route entry 9999::/64 to Switch B, then check the BGP route table on Switch B

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2                9.9.9.9            9 10 i
admin@XorPlus# run show bgp routes ipv6 deatil
Invalid IPNet
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 9 10
    Nexthop: 3003::2
    Multiple Exit Discriminator: 0
    Local Preference: 100
admin@XorPlus#
```

Step 3: Enable BGP Route Dampening on Switch B

```
admin@XorPlus#set protocols bgp damping disable false
```

Note: The default half-life: 15, the default max-suppress: 60, the default reuse: 750, the default suppress: 3000.

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2                9.9.9.9            9 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
    From peer: 9.9.9.9
    Route: Winner
    Origin: IGP
    AS Path: 9 10
    Nexthop: 3003::2
    Multiple Exit Discriminator: 0
    Local Preference: 100
admin@XorPlus#
```

Step 4: Withdraw and update the BGP route entry on Switch A and repeat the operation three times, then check the BGP route table, the BGP route entry should be suppressed after the third times.

The first time:

Switch B BGP route table:

```
admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
```

```

Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2            9.9.9.9            9 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 3003::2
  Multiple Exit Discriminator: 0
  Local Preference: 100
admin@XorPlus#

```

The second time:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2            9.9.9.9            9 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 3003::2
  Multiple Exit Discriminator: 0
  Local Preference: 100
admin@XorPlus#

```

The third time:

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
admin@XorPlus#

```

Note: We can see that the BGP route was suppressed on Switch B the third time you've withdrawn the BGP route entry, and the BGP route entry will be recover after 60 minutes.

60 minutes later :

```

admin@XorPlus# run show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete
  Prefix                Nexthop                Peer                AS Path
  -----                -
*> 9999::/64            3003::2            9.9.9.9            9 10 i
admin@XorPlus# run show bgp routes ipv6 detail
9999::/64
  From peer: 9.9.9.9
  Route: Winner
  Origin: IGP
  AS Path: 9 10
  Nexthop: 3003::2

```

```
Multiple Exit Discriminator: 0
Local Preference: 100
admin@XorPlus#
```

4.2 BGP fast-external-falover

BGP fast-external-falover Commands References:

```
admin@XorPlus#set protocols bgp fast-external-falover disable [true/false]
```

Note: BGP fast-external-falover is enabled by default. The BGP fast-external-falover command is used to disable or enable fast-external-falover for BGP peering sessions with directly connected external peers. The session is immediately reset if link goes down, and only directly connected peering sessions are supported.

If BGP fast-external-falover is disabled, The BGP routing process will wait until the default hold timer expires to reset the peering session.

Configuration Example 2:

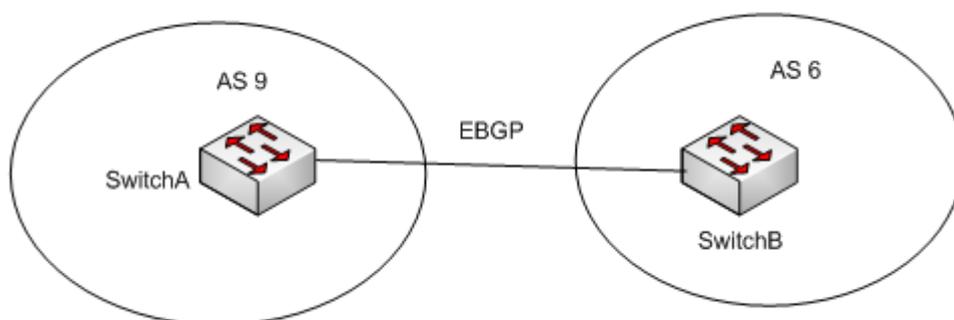


Figure 4-3

Step 1: SwitchA, SwitchB Establishing EBGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 9.9.9.9
admin@XorPlus# set protocols bgp local-as "9"
admin@XorPlus# set protocols bgp peer 3003::1 local-ip "3003::2"
admin@XorPlus# set protocols bgp peer 3003::1 as "6"
admin@XorPlus# set protocols bgp peer 3003::1 ipv6-unicast true
```

Switch B:

```
admin@XorPlus# set protocols bgp bgp-id 6.6.6.6
admin@XorPlus# set protocols bgp local-as "6"
admin@XorPlus# set protocols bgp peer 3003::2 local-ip "3003::1"
admin@XorPlus# set protocols bgp peer 3003::2 as "9"
admin@XorPlus# set protocols bgp peer 3003::2 ipv6-unicast true
```

Then check BGP peer status on Switch A, SwitchB:

Switch A:

```

admin@XorPlus# run show bgp peers detail
Peer 1: local 3003::2/179 remote 3003::1/38709
Peer ID: 6.6.6.6
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 6
Updates Received: 0, Updates Sent: 7
Messages Received: 10, Messages Sent: 17
Time since last received update: n/a
Number of transitions to ESTABLISHED: 2
Time since last entering ESTABLISHED state: 230 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Switch B:

```

admin@XorPlus# run show bgp peers detail
Peer 2: local 3003::1/38709 remote 3003::2/179
Peer ID: 9.9.9.9
Peer State: ESTABLISHED
Admin State: START
Negotiated BGP Version: 4
Peer AS Number: 9
Updates Received: 7, Updates Sent: 0
Messages Received: 18, Messages Sent: 11
Time since last received update: 227 seconds
Number of transitions to ESTABLISHED: 1
Time since last entering ESTABLISHED state: 255 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds

```

Step 2: Shutdown the link BGP peer connected then check the BGP peer status, the BGP peer will be down immediately as BGP fast-external-falover is enable by default.

```

admin@XorPlus# set interface gigabit-ethernet te-1/1/46 disable true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/179
Peer ID: none
Peer State: CONNECT
Admin State: START
Negotiated BGP Version: n/a
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 0, Messages Sent: 0
Time since last received update: n/a
Number of transitions to ESTABLISHED: 1
Time since last in ESTABLISHED state: 11 seconds
Retry Interval: 120 seconds
Hold Time: n/a, Keep Alive Time: n/a
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds

```

```
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#
```

Step 3: Up the link BGP peer connected and disable BGP fast-external-falover on Switch A

```
admin@XorPlus# set interface gigabit-ethernet te-1/1/46 disable false
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols bgp fast-external-falover disable true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/52347 remote 3003::2/179
  Peer ID: 9.9.9.9
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 9
  Updates Received: 1, Updates Sent: 0
  Messages Received: 3, Messages Sent: 2
  Time since last received update: 4 seconds
  Number of transitions to ESTABLISHED: 2
  Time since last entering ESTABLISHED state: 4 seconds
  Retry Interval: 120 seconds
  Hold Time: 90 seconds, Keep Alive Time: 30 seconds
  Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
  Minimum AS Origination Interval: 0 seconds
  Minimum Route Advertisement Interval: 0 seconds
```

Step 4: Shutdown the link BGP peer connected, then check the BGP peer status, the BGP peer will not shutdown immediately, but it will goe down after 90 seconds

```
admin@XorPlus# set interface gigabit-ethernet te-1/1/46 disable true
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/52347 remote 3003::2/179
  Peer ID: 9.9.9.9
  Peer State: ESTABLISHED
  Admin State: START
  Negotiated BGP Version: 4
  Peer AS Number: 9
  Updates Received: 1, Updates Sent: 0
  Messages Received: 7, Messages Sent: 7
  Time since last received update: 139 seconds
  Number of transitions to ESTABLISHED: 2
  Time since last entering ESTABLISHED state: 139 seconds
  Retry Interval: 120 seconds
  Hold Time: 90 seconds, Keep Alive Time: 30 seconds
  Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
  Minimum AS Origination Interval: 0 seconds
  Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#
```

90 seconds later :

```
admin@XorPlus# run show bgp peers detail 3003::2
Peer 1: local 3003::1/179 remote 3003::2/179
Peer ID: none
Peer State: ACTIVE
Admin State: START
Negotiated BGP Version: n/a
Peer AS Number: 9
Updates Received: 0, Updates Sent: 0
Messages Received: 0, Messages Sent: 0
Time since last received update: n/a
Number of transitions to ESTABLISHED: 2
Time since last in ESTABLISHED state: 34 seconds
Retry Interval: 120 seconds
Hold Time: n/a, Keep Alive Time: n/a
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
admin@XorPlus#
```

Label BGP

This feature module describes how to add label mapping information to the Border Gateway Protocol (BGP) message that is used to distribute the route on Pica8 Switch.

- Labeled BGP Support
- Configuration Example for Labeled Support
 - IPV4 Labeled BGP Configuration
 - IPV6 Labeled-BGP Configuration
- Debugging CLI for Labeled-BGP

Labeled BGP Support

Restrictions

- The Pica8 Switch supports only the client functionality of RFC 3107 and not its area border router (ABR) functionality.
- The Pica8 Switch does not support two label-pop (Label pop is the process of removing label header).
- MPLS label push is not supported.
- MPLS Ping and MPLS traceroute is not supported.

Overview

The Labeled BGP Support feature provides the option to use the BGP update message (that is used to distribute the route) to re-distribute Multiprotocol Label Switching (MPLS) label mapped to that route. The label mapping information is added (using send-label option of RFC 3107) to the same BGP message that is used to distribute the route. This process is useful in inter-domain routing.

How to Configure

Procedure

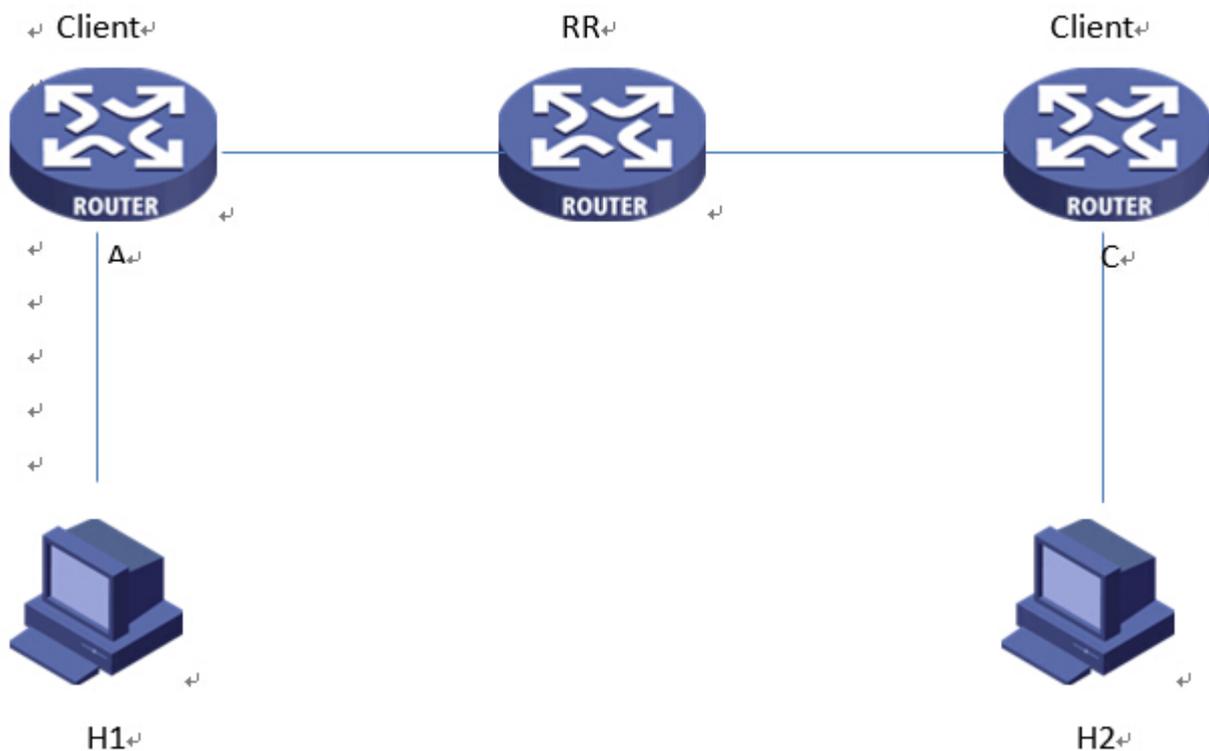
	Command or Action	Purpose
Step1	configure terminal Example: admin@XorPlus> configure	Enters the global configuration mode.
Step2	set protocols bgp bgp-id <IPv4> Example: admin@XorPlus# set protocols bgp bgp-id 1.1.1.1	Configure BGP router-id
Step3	set protocols bgp local-as <text> Example: admin@XorPlus# set protocols bgp local-as 1	Configure bgp AS Number, it has two types: <1..4294967295> or 4-byte-as-number <0..65535>.<0..65535>
Step4	set protocols bgp peer <peer ip address> as <as-number> Example: admin@XorPlus# set protocols bgp peer 1.1.1.1 as 1	Configure BGP peer and specify the as number of the BGP peer
Step5	set protocols bgp peer <peer ip address> local-ip <ip-address> Example: admin@XorPlus# set protocols bgp peer 1.1.1.1 local-ip 2.2.2.2	Configure BGP peer and specify the local-IP address for this BGP peer
Step6	set protocols bgp peer <peer ip address> ipv4-label-unicast/ ipv6-label-unicast <true/false> Example: admin@XorPlus# set protocols bgp peer 1.1.1.1 ipv4-label-unicast true	Configures the Pica8 Switch to associate a BGP label to the prefix
Step7	set protocols mpls in-label [16..1048575] connect/connect6 <IPv4Net>/<IPv6Net> Example: admin@XorPlus# set protocols mpls in-label 44 connect 192.168.2.0/24	Configure the MPLS label mapping for the top level label, then the destination network will associate with the top level mpls label
Step8	set protocols mpls in-label [16..1048575] next-hop <IPv4>/ <IPv6> Example: admin@XorPlus# set protocols mpls in-label 44 next-hop 1.1.1.1	Configure the MPLS label mapping where MPLS packet out of border BGP router, the next-hop indicates where the MPLS packet will be forwarded and with a label you configured

Configuration Example for Labeled Support

- IPV4 Labeled BGP Configuration
- IPV6 Labeled-BGP Configuration

IPV4 Labeled BGP Configuration

Scenario 1: Incoming Three Layer Label



Step 1: Configure loopback interface

Switch A:

```
admin@XorPlus#set vlan-interface loopback address 28.28.28.28 prefix-length 32
```

Switch B:

```
admin@XorPlus#set vlan-interface loopback address 17.17.17.17 prefix-length 32
```

Switch C:

```
admin@XorPlus#set vlan-interface loopback address 31.31.31.31 prefix-length 32
```

Step 2: Configure VLAN-interface

Switch A:

```
admin@XorPlus# set vlans vlan-id 200 13-interface "vlan200"
admin@XorPlus#set vlans vlan-id 300 13-interface "vlan300"
admin@XorPlus# set vlan-interface interface vlan200 vif vlan200 address 192.168.2.1
```

```

prefix-length 24
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 192.168.3.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus#set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200

```

Switch B:

```

admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.2
prefix-length 24
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 192.168.5.2
prefix-length 24
admin@XorPlus#set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300

```

Switch C:

```

admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 192.168.5.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan700 vif vlan700 address 192.168.7.1
prefix-length 24
admin@XorPlus#set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700

```

Step 3: Configure static route on A, B, C so that the loopback interface is reachable.**Switch A:**

```

admin@XorPlus#set protocols static route 17.17.17.17/32 next-hop 192.168.3.2

```

Switch B:

```

admin@XorPlus#set protocols static route 28.28.28.28/32 next-hop 192.168.3.1
admin@XorPlus#set protocols static route 31.31.31.31/32 next-hop 192.168.5.1

```

Switch C:

```

admin@XorPlus#set protocols static route 17.17.17.17/32 next-hop 192.168.5.2

```

Step 4: Configure BGP**Switch A:**

```

admin@XorPlus#set protocols bgp bgp-id 28.28.28.28
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp fast-external-fallover disable true
admin@XorPlus#set protocols bgp peer 17.17.17.17 local-ip "28.28.28.28"

```

```
admin@XorPlus#set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus#set protocols bgp peer 17.17.17.17 next-hop-self true
```

Switch B:

```
admin@XorPlus#set protocols bgp bgp-id 17.17.17.17
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus#set protocols bgp peer 28.28.28.28 local-ip "17.17.17.17"
admin@XorPlus#set protocols bgp peer 28.28.28.28 as "17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 next-hop-self true
admin@XorPlus#set protocols bgp peer 28.28.28.28 client true
admin@XorPlus#set protocols bgp peer 31.31.31.31 local-ip "17.17.17.17"
admin@XorPlus#set protocols bgp peer 31.31.31.31 as "17"
admin@XorPlus#set protocols bgp peer 31.31.31.31 next-hop-self true
admin@XorPlus#set protocols bgp peer 31.31.31.31 client true
```

Switch C:

```
admin@XorPlus#set protocols bgp bgp-id 31.31.31.31
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp peer 17.17.17.17 local-ip "31.31.31.31"
admin@XorPlus#set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus#set protocols bgp peer 17.17.17.17 next-hop-self true
```

Step 5: Enable Labeled-BGP on Switch A, Switch B, Switch C**Switch A:**

```
admin@XorPlus#set protocols bgp peer 17.17.17.17 ipv4-label-unicast true
```

Switch B:

```
admin@XorPlus#set protocols bgp peer 28.28.28.28 ipv4-label-unicast true
admin@XorPlus# set protocols bgp peer 31.31.31.31 ipv4-label-unicast true
```

Switch C:

```
admin@XorPlus#set protocols bgp peer 17.17.17.17 ipv4-label-unicast true
```

Step 6: Distribute a BGP route**Switch A:**

```
admin@XorPlus#set protocols bgp network4 192.168.2.0/24
```

Step 7: Check BGP route table**Switch A:**

```
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop                In-label/Out-label
```

```

-----
*> 192.168.2.0/30      0.0.0.0      nolabel/nolabel
admin@XorPlus#

```

Switch B:

```

admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop          In-label/Out-label
  -----
*> 192.168.2.0/30      28.28.28.28     16/imp-null
admin@XorPlus#

```

Switch C:

```

admin@XorPlus# run show bgp routes ipv4 l
Possible completions:
  <IPNet>                Print BGP IPv4 routes of specified prefix
  label                  Print BGP IPv4 routes with labels
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop          In-label/Out-label
  -----
*> 192.168.2.0/30      17.17.17.17     17/16
admin@XorPlus#

```

Step 8: Check MPLS forwarding table on Switch A, Switch B, Switch C**Switch A:**

```

admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
admin@XorPlus#

```

Switch B:

```

admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
16        imp-null   192.168.2.0/30  te-1/1/2            192.168.3.1
admin@XorPlus#

```

Switch C:

```

admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
17        16        192.168.2.0/30  te-1/1/44          192.168.5.2
admin@XorPlus#

```

Step 9: Configure MPLS label mapping on Switch C**Switch C:**

```
admin@XorPlus#set protocols mpls in-label 1000 connect 192.168.2.0/30
```

Note : The label 1000 is the top label that is incoming to the BGP domain.

Step 10: Configure MPLS label mapping on Switch A

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop 192.168.2.2
```

Note: The label 2000 is the second-level label that is incoming to the BGP domain, and it's the outgoing label that is out of BGP domain. The next-hop indicates the next-hop IP address where the MPLS packet is forwarding, and with a label 2000.

Step 11: Check the MPLS forwarding table on Switch A, Switch B, and Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
2000      2000        --              te-1/1/38           192.168.2.2
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
16        imp-null   192.168.2.0/30  te-1/1/2            192.168.3.1
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
17         16         192.168.2.0/30  te-1/1/44           192.168.5.2
1000      16         192.168.2.0/30  te-1/1/44           192.168.5.2
admin@XorPlus#
```

Step 12: H2 sends MPLS packet with 3-level MPLS label.

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 00 03 00 00 00
MAC: Type                = 0x8847 (Ethernet II)
MAC:
MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 1000 (0x3E8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live       = 255 (0xFF)
MPLS: Label 2
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live       = 100 (0x64)
MPLS: Label 3
MPLS: Label              = 3000 (0xBB8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live       = 0 (0x0)
MPLS:
IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length        = 20 (0x14)
IP: Type of Service      = 00 (0x00)
IP: 000.....            = Routine
IP: ...0....            = Normal Delay
IP: ....0...            = Normal Throughput
IP: .....0..           = Normal Reliability
IP: .....0.            = Normal Cost
IP: .....0             = Reserved
IP: Total Length         = 226 (0x00E2)
IP: Identification      = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset      = 0
IP: Time to Live         = 0 (0x00)
IP: Protocol             = 255 (0xFF)
IP: Checksum             = 0x2EC9
IP: Source Address       = 192.168.7.2
IP: Destination Address  = 192.168.2.2
IP:

```

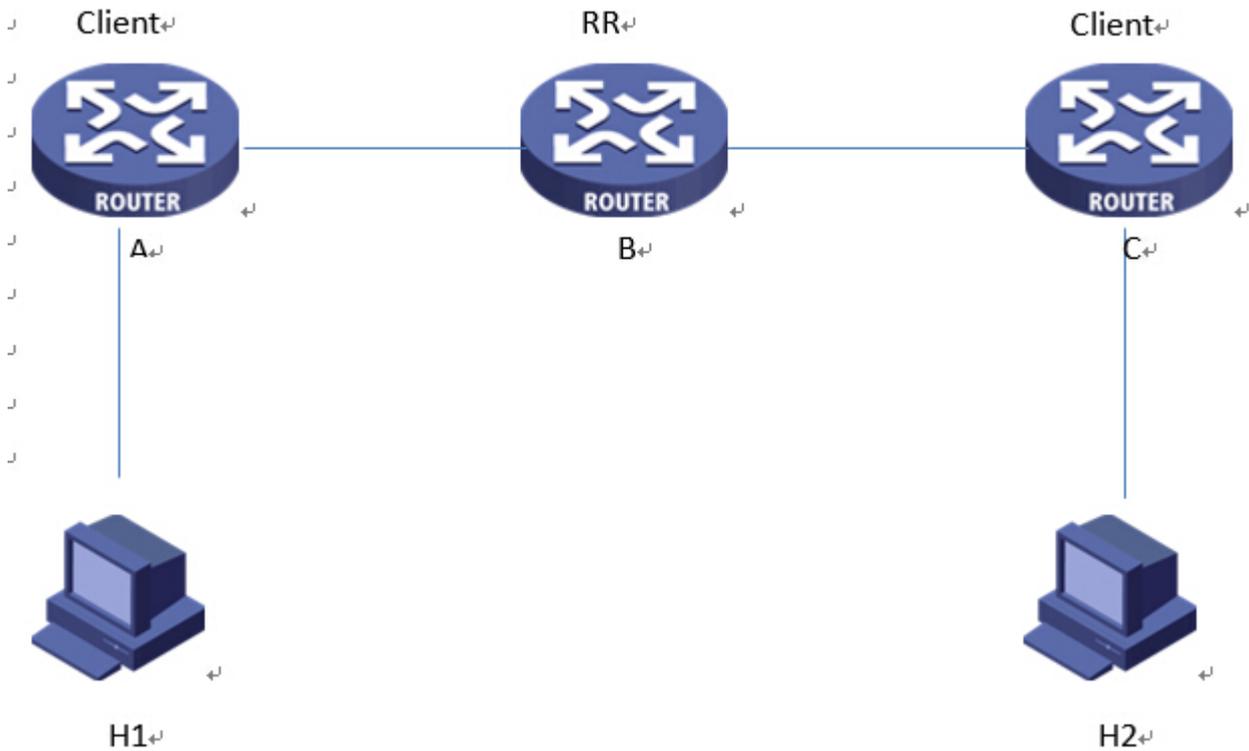
Step 13: Check the packet out of Switch A. the top label 1000 will be popped at the Penultimate Hop.

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 22 22 22 22 22
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                = 0x8847 (Ethernet II)
MAC:
MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live       = 252 (0xFC)
MPLS: Label 2
MPLS: Label              = 3000 (0xBB8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live       = 0 (0x0)
MPLS:
IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length        = 20 (0x14)
IP: Type of Service      = 00 (0x00)
IP: 000.....            = Routine
IP: ...0....            = Normal Delay
IP: ....0...            = Normal Throughput
IP: .....0..           = Normal Reliability
IP: .....0.            = Normal Cost
IP: .....0             = Reserved
IP: Total Length         = 226 (0x00E2)
IP: Identification      = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset      = 0
IP: Time to Live         = 0 (0x00)
IP: Protocol             = 255 (0xFF)
IP: Checksum             = 0x2EC9
IP: Source Address       = 192.168.7.2
IP: Destination Address  = 192.168.2.2
IP:

```

Scenario 2: Incoming Two Layer Label



Step 1: Configure loopback interface

Switch A:

```
admin@XorPlus# set vlan-interface loopback address 28.28.28.28 prefix-length 32
```

Switch B:

```
admin@XorPlus# set vlan-interface loopback address 17.17.17.17 prefix-length 32
```

Switch C:

```
admin@XorPlus# set vlan-interface loopback address 31.31.31.31 prefix-length 32
```

Step 2: Configure VLAN-interface

Switch A:

```
admin@XorPlus# set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlan-interface interface vlan200 vif vlan200 address 192.168.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

Switch B:

```
admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.2
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300
```

Switch C:

```
admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan700 vif vlan700 address 192.168.7.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700
```

Step 3: Configure static route on A,B,C so that the loopback interface is reachable**Switch A:**

```
admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.3.2
```

Switch B:

```
admin@XorPlus# set protocols static route 28.28.28.28/32 next-hop 192.168.3.1
admin@XorPlus# set protocols static route 31.31.31.31/32 next-hop 192.168.5.1
```

Switch C:

```
admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.5.2
```

Step 4: Configure BGP**Switch A:**

```
admin@XorPlus# set protocols bgp bgp-id 28.28.28.28
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp fast-external-fallover disable true
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "28.28.28.28"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true
```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 17.17.17.17
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus# set protocols bgp peer 28.28.28.28 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 as "17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 next-hop-self true
admin@XorPlus# set protocols bgp peer 28.28.28.28 client true
admin@XorPlus# set protocols bgp peer 31.31.31.31 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 as "17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 next-hop-self true
admin@XorPlus# set protocols bgp peer 31.31.31.31 client true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 31.31.31.31
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "31.31.31.31"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true

```

Step 5: Enable Labeled-BGP on switch A,B,C

Switch A:

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp peer 28.28.28.28 ipv4-label-unicast true
admin@XorPlus# set protocols bgp peer 31.31.31.31 ipv4-label-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Step 6: Switch A distributes a BGP route

Switch A:

```

admin@XorPlus# set protocols bgp network4 192.168.2.0/24

```

Step 7: Check BGP route table

Switch A:

```

admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop                In-label/Out-label
  -----                -
*> 192.168.2.0/30      0.0.0.0                nolabel/nolabel
admin@XorPlus#

```

Switch B:

```
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	28.28.28.28	16/imp-null

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show bgp routes ipv4 l
Possible completions:
<IPNet>          Print BGP IPv4 routes of specified prefix
label           Print BGP IPv4 routes with labels
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	17.17.17.17	17/16

```
admin@XorPlus#
```

Step 8: Check MPLS forwarding table on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
16        imp-null   192.168.2.0/30   te-1/1/2            192.168.3.1
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----
17        16        192.168.2.0/30   te-1/1/44          192.168.5.2
admin@XorPlus#
```

Step 9: Configure MPLS label mapping on Switch C

Switch C:

```
admin@XorPlus# set protocols mpls in-label 1000 connect 192.168.2.0/30
```

Note : The label 1000 is the top label that is incoming to the BGP domain

Step 10: Configure MPLS label mapping on Switch A

```
admin@XorPlus# set protocols mpls in-label 2000 next-hop 192.168.2.2
```

Note : The label 2000 is the second-level label that incoming BGP domain , and it's the outgoing label that out of BGP domain the next-hop is indicate the next-hop ip address where the mpls packet forwarding and with a label 2000.

Step 11: Check the MPLS forwarding table on Switch A, S witch B, S witch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
2000      2000        --              te-1/1/38           192.168.2.2
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
16        imp-null   192.168.2.0/30  te-1/1/2            192.168.3.1
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
17        16         192.168.2.0/30  te-1/1/44           192.168.5.2
1000     16         192.168.2.0/30  te-1/1/44           192.168.5.2
admin@XorPlus#
```

Step 12: H2 send MPLS packet with 2-level MPLS label

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 00 03 00 00 00
MAC: Type                = 0x8847 (Ethernet II)
MAC:
MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 1000 (0x3E8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live      = 255 (0xFF)
MPLS: Label 2
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live      = 100 (0x64)
MPLS:
IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    ....0... = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 230 (0x00E6)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 0 (0x00)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0x2EC5
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:

```

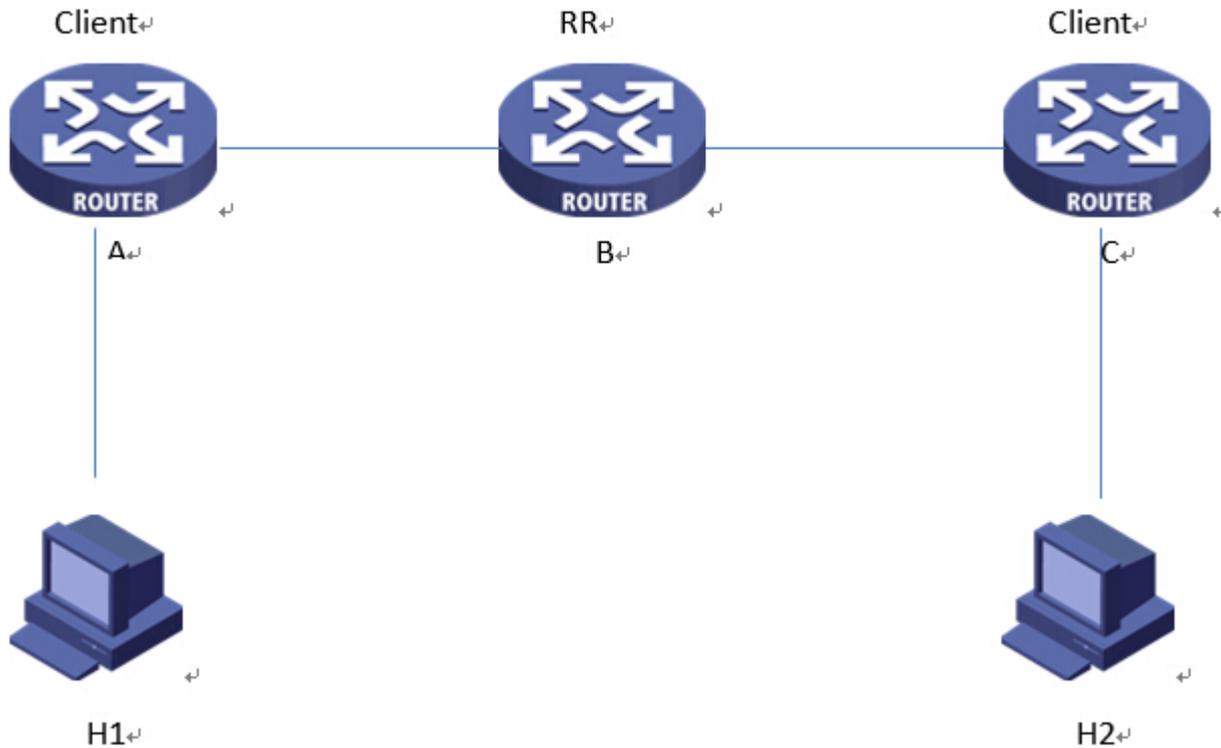
Step 13: Check the packet out of Switch A, the top label 1000 will be popped at the Penultimate Hop

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 22 22 22 22 22
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                = 0x8847 (Ethernet II)
MAC:
MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live      = 252 (0xFC)
MPLS:
IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    ....0... = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 230 (0x00E6)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 0 (0x00)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0x2EC5
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:
IP: TransportLayer : Undecoded Protocol (255)

```

Scenario 3: Incoming One Layer Label



Step 1: Configure Loopback interface

Switch A:

```
admin@XorPlus# set vlan-interface loopback address 28.28.28.28 prefix-length 32
```

Switch B:

```
admin@XorPlus# set vlan-interface loopback address 17.17.17.17 prefix-length 32
```

Switch C:

```
admin@XorPlus# set vlan-interface loopback address 31.31.31.31 prefix-length 32
```

Step 2: Configure VLAN-interface

Switch A:

```
admin@XorPlus# set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlan-interface interface vlan200 vif vlan200 address 192.168.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

Switch B:

```
admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.2
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300
```

Switch C:

```
admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan700 vif vlan700 address 192.168.7.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700
```

Step 3: Configure static route on Switch A, Switch B, Switch C so that the loopback interface is reachable

Switch A:

```
admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.3.2
```

Switch B:

```
admin@XorPlus# set protocols static route 28.28.28.28/32 next-hop 192.168.3.1
admin@XorPlus# set protocols static route 31.31.31.31/32 next-hop 192.168.5.1
```

Switch C:

```
admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.5.2
```

Step 4: Configure BGP

Switch A:

```
admin@XorPlus# set protocols bgp bgp-id 28.28.28.28
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp fast-external-fallover disable true
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "28.28.28.28"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true
```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 17.17.17.17
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus# set protocols bgp peer 28.28.28.28 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 as "17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 next-hop-self true
admin@XorPlus# set protocols bgp peer 28.28.28.28 client true
admin@XorPlus# set protocols bgp peer 31.31.31.31 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 as "17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 next-hop-self true
admin@XorPlus# set protocols bgp peer 31.31.31.31 client true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 31.31.31.31
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "31.31.31.31"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true

```

Step 5: Enable Labeled-BGP on Switch A, Switch B, Switch C

Switch A:

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp peer 28.28.28.28 ipv4-label-unicast true
admin@XorPlus# set protocols bgp peer 31.31.31.31 ipv4-label-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Step 6: Switch A distribute a BGP route

Switch A:

```

admin@XorPlus# set protocols bgp network4 192.168.2.0/24

```

Step 7: Check BGP route table

Switch A:

```

admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop                In-label/Out-label
  -----                -
*> 192.168.2.0/30      0.0.0.0                nolabel/nolabel
admin@XorPlus#

```

Switch B:

```
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	28.28.28.28	16/imp-null

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show bgp routes ipv4 l
```

Possible completions:

```
<IPNet>          Print BGP IPv4 routes of specified prefix
label           Print BGP IPv4 routes with labels
```

```
admin@XorPlus# run show bgp routes ipv4 label
```

```
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	17.17.17.17	17/16

```
admin@XorPlus#
```

Step 8: Check MPLS forwarding table on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
admin@XorPlus#				

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
16	imp-null	192.168.2.0/30	te-1/1/2	192.168.3.1

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
17	16	192.168.2.0/30	te-1/1/44	192.168.5.2

```
admin@XorPlus#
```

Step 9: Configure MPLS label mapping on Switch C

Switch C:

```
admin@XorPlus# set protocols mpls in-label 1000 connect 192.168.2.0/30
```

Note: The label 1000 is the top label that incoming BGP domain

Step 10: Configure MPLS label mapping on Switch A

```
admin@XorPlus# set protocols mpls in-label 2000 next-hop 192.168.2.2
```

Note: The label 2000 is the second-level label that is incoming to the BGP domain, and it's the outgoing label that is out of BGP domain, the next-hop is indicating the next-hop IP address where the MPLS packet is forwarding and with a label 2000.

Step 11: Check the MPLS forwarding table on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
2000      2000        --              te-1/1/38           192.168.2.2
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
16        imp-null   192.168.2.0/30  te-1/1/2            192.168.3.1
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
17        16         192.168.2.0/30  te-1/1/44           192.168.5.2
1000     16         192.168.2.0/30  te-1/1/44           192.168.5.2
admin@XorPlus#
```

Step 12: H2 send MPLS packet with 1-level MPLS label

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 00 03 00 00 00
MAC: Type                = 0x8847 (Ethernet II)
MAC:

MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 1000 (0x3E8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live      = 4 (0x4)
MPLS:

IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    .....0.. = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 234 (0x00EA)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 255 (0xFF)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0x2FC0
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:
IP: TransportLayer : Undecoded Protocol (255)

```

Step 13: Check the packet out of Switch A, the top label 1000 will be popped at the Penultimate Hop

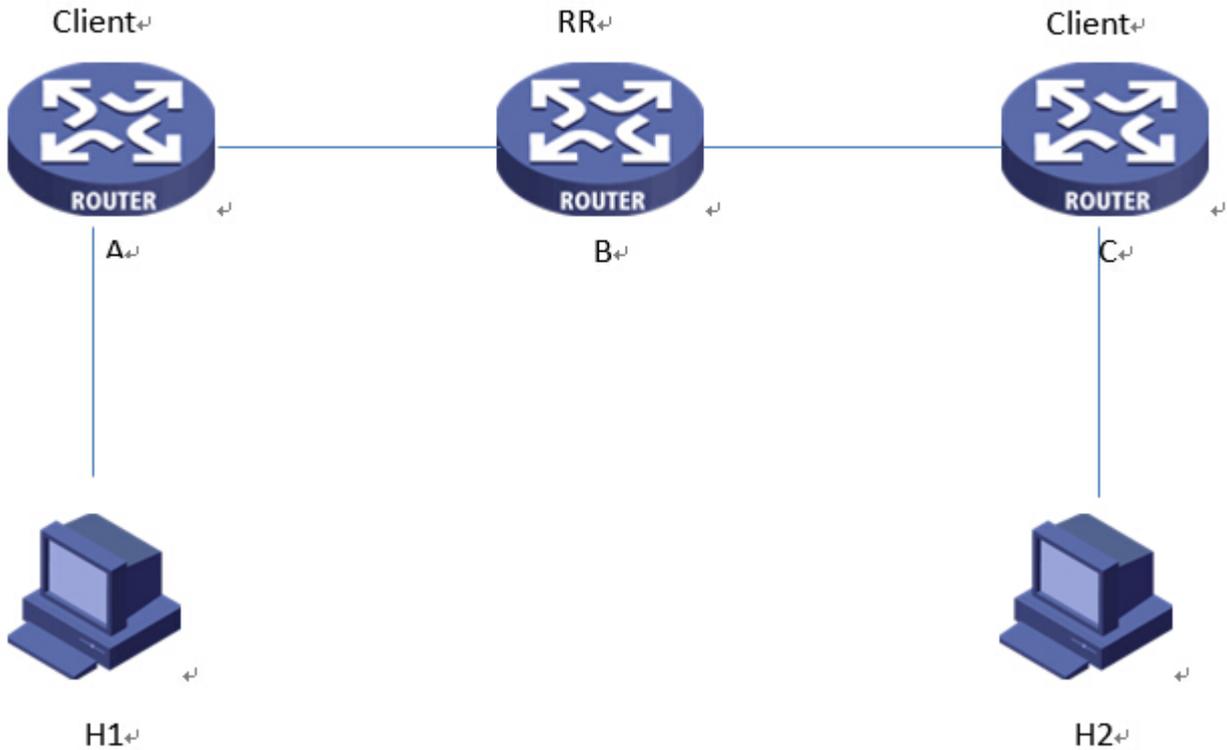
```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 22 22 22 22 22
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                : 0x0800 (Ethernet II)
MAC:

IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    .....0.. = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 234 (0x00EA)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 1 (0x01)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0x2DC1
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:
IP: TransportLayer : Undecoded Protocol (255)

```

Scenario 4: Incoming IP Packet with No-label



Step 1: Configure Loopback interface

Switch A:

```
admin@XorPlus# set vlan-interface loopback address 28.28.28.28 prefix-length 32
```

Switch B:

```
admin@XorPlus# set vlan-interface loopback address 17.17.17.17 prefix-length 32
```

Switch C:

```
admin@XorPlus# set vlan-interface loopback address 31.31.31.31 prefix-length 32
```

Step 2: Configure VLAN-interface

Switch A:

```
admin@XorPlus# set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlan-interface interface vlan200 vif vlan200 address 192.168.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus# set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

Switch B:

```

admin@XorPlus# set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlan-interface interface vlan300 vif vlan300 address 192.168.3.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.2
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300

```

Switch C:

```

admin@XorPlus# set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus# set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus# set vlan-interface interface vlan500 vif vlan500 address 192.168.5.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan700 vif vlan700 address 192.168.7.1
prefix-length 24
admin@XorPlus# set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus# set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700

```

Step 3: Configure static route on Switch A, Switch B, Switch C so that the loopback interface is reachable

Switch A:

```

admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.3.2

```

Switch B:

```

admin@XorPlus# set protocols static route 28.28.28.28/32 next-hop 192.168.3.1
admin@XorPlus# set protocols static route 31.31.31.31/32 next-hop 192.168.5.1

```

Switch C:

```

admin@XorPlus# set protocols static route 17.17.17.17/32 next-hop 192.168.5.2

```

Step 4: Configure BGP

Switch A:

```

admin@XorPlus# set protocols bgp bgp-id 28.28.28.28
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp fast-external-fallover disable true
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "28.28.28.28"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true

```

Switch B:

```

admin@XorPlus# set protocols bgp bgp-id 17.17.17.17
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus# set protocols bgp peer 28.28.28.28 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 as "17"
admin@XorPlus# set protocols bgp peer 28.28.28.28 next-hop-self true
admin@XorPlus# set protocols bgp peer 28.28.28.28 client true
admin@XorPlus# set protocols bgp peer 31.31.31.31 local-ip "17.17.17.17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 as "17"
admin@XorPlus# set protocols bgp peer 31.31.31.31 next-hop-self true
admin@XorPlus# set protocols bgp peer 31.31.31.31 client true

```

Switch C:

```

admin@XorPlus# set protocols bgp bgp-id 31.31.31.31
admin@XorPlus# set protocols bgp local-as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 local-ip "31.31.31.31"
admin@XorPlus# set protocols bgp peer 17.17.17.17 as "17"
admin@XorPlus# set protocols bgp peer 17.17.17.17 next-hop-self true

```

Step 5: Enable Labeled-BGP on Switch A, Switch B, Switch C**Switch A:**

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Switch B:

```

admin@XorPlus# set protocols bgp peer 28.28.28.28 ipv4-label-unicast true
admin@XorPlus# set protocols bgp peer 31.31.31.31 ipv4-label-unicast true

```

Switch C:

```

admin@XorPlus# set protocols bgp peer 17.17.17.17 ipv4-label-unicast true

```

Step 6: Switch A distribute a BGP route**Switch A:**

```

admin@XorPlus# set protocols bgp network4 192.168.2.0/24

```

Step 7: Check BGP route table

```

switchA:
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route

  Prefix                Nexthop                In-label/Out-label
  -----
*> 192.168.2.0/30      0.0.0.0                no-label/no-label
admin@XorPlus#

```

Switch B:

```
admin@XorPlus# run show bgp routes ipv4 label
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	28.28.28.28	16/imp-null

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show bgp routes ipv4 l
```

Possible completions:

```
<IPNet>          Print BGP IPv4 routes of specified prefix
label           Print BGP IPv4 routes with labels
```

```
admin@XorPlus# run show bgp routes ipv4 label
```

```
Status Codes: * valid route, > best route
```

Prefix	Nexthop	In-label/Out-label
*> 192.168.2.0/30	17.17.17.17	17/16

```
admin@XorPlus#
```

Step 8: Check MPLS forwarding table on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
-----	-----	-----	-----	-----

```
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
-----	-----	-----	-----	-----
16	imp-null	192.168.2.0/30	te-1/1/2	192.168.3.1

```
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
```

In-label	Out-label	Destination	Outgoing-Interface	Next Hop
-----	-----	-----	-----	-----
17	16	192.168.2.0/30	te-1/1/44	192.168.5.2

```
admin@XorPlus#
```

Step 9: Configure MPLS label mapping on Switch C

Switch C:

```
admin@XorPlus# set protocols mpls in-label 1000 connect 192.168.2.0/30
```

Note : The label 1000 is the top label that is incoming to the BGP domain

Step 10: Configure MPLS label mapping on Switch A

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop 192.168.2.2
```

Note : The label 2000 is the second-level label that is incoming to the BGP domain, and it's the outgoing label that comes out of BGP domain, the next-hop indicates the next-hop IP address where the MPLS packet forwards with a label 2000.

Step 11: Check the MPLS forwarding table on Switch A, Switch B, Switch C

Switch A:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
2000      2000        --              te-1/1/38           192.168.2.2
admin@XorPlus#
```

Switch B:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
16        imp-null   192.168.2.0/30   te-1/1/2            192.168.3.1
admin@XorPlus#
```

Switch C:

```
admin@XorPlus# run show mpls forwarding-table ipv4 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
17        16         192.168.2.0/30   te-1/1/44           192.168.5.2
1000     16         192.168.2.0/30   te-1/1/44           192.168.5.2
admin@XorPlus#
```

Step 12: H2 send IP packet

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 00 03 00 00 00
MAC: Type                : 0x0800 (Ethernet II)
MAC:

IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    ....0..  = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 238 (0x00EE)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 64 (0x40)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0xEEBC
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:
IP: TransportLayer : Undecoded Protocol (255)

```

Step 13: Check the packet out of Switch A

```

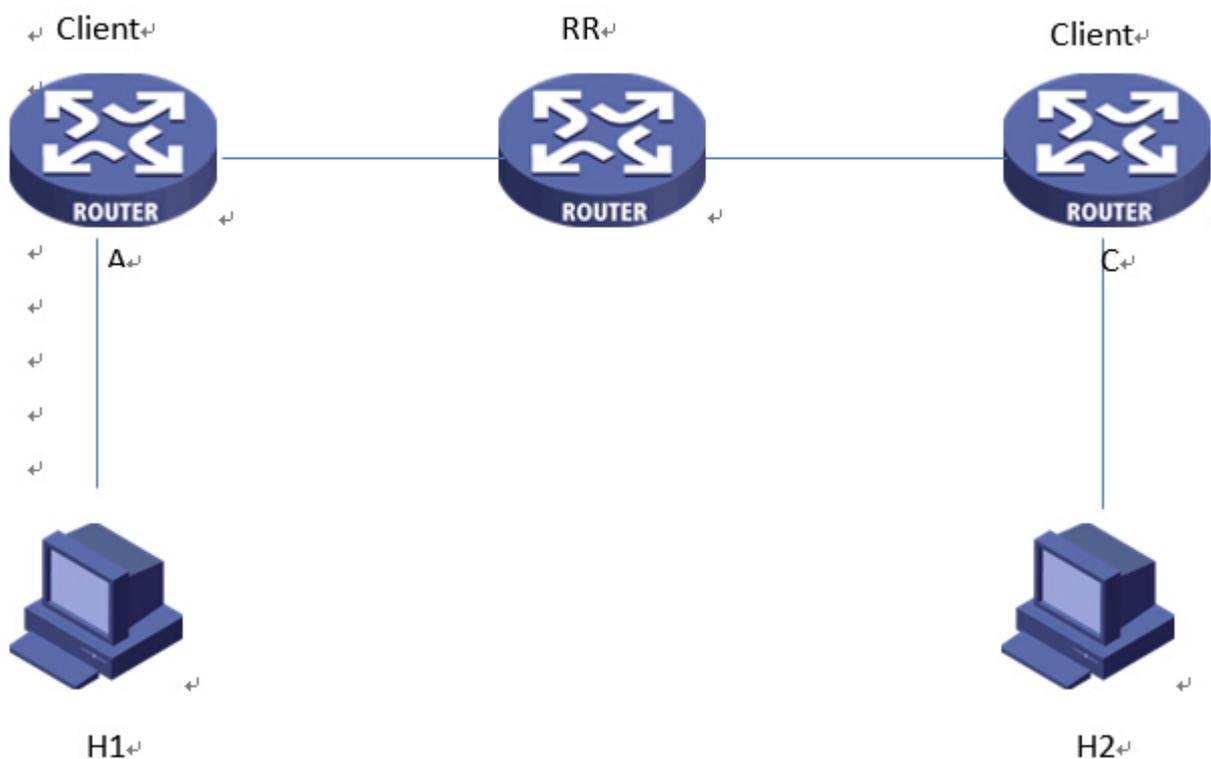
MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 22 22 22 22 22
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                : 0x0800 (Ethernet II)
MAC:

IP: ----- IP Header -----
IP:
IP: Version              = 04 (0x04)
IP: Header Length       = 20 (0x14)
IP: Type of Service     = 00 (0x00)
IP:                    000..... = Routine
IP:                    ...0.... = Normal Delay
IP:                    ...0...  = Normal Throughput
IP:                    ....0..  = Normal Reliability
IP:                    .....0.  = Normal Cost
IP:                    .....0   = Reserved
IP: Total Length        = 238 (0x00EE)
IP: Identification     = 0 (0x0000)
IP: Flags Bit1 .0. May Fragment
IP: Flags Bit2 ..0 Last Fragment
IP: Fragment Offset    = 0
IP: Time to Live       = 61 (0x3D)
IP: Protocol           = 255 (0xFF)
IP: Checksum           = 0xF1BC
IP: Source Address     = 192.168.7.2
IP: Destination Address = 192.168.2.2
IP:
IP: TransportLayer : Undecoded Protocol (255)

```

IPv6 Labeled-BGP Configuration

Scenario 1: Incoming three layer label



Step 1: Configure Loopback interface

switchA:

```
admin@XorPlus#set vlan-interface loopback address 2828::28 prefix-length 128
```

switchB:

```
admin@XorPlus#set vlan-interface loopback address 1717::17 prefix-length 128
```

switchC:

```
admin@XorPlus#set vlan-interface loopback address 3131::31 prefix-length 128
```

Step 2: Configure Vlan-interface

switchA:

```
admin@XorPlus#set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlan-interface interface vlan200 vif vlan200 address 2001::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus#set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

switchB:

```
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::2
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::2
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300
```

switchC:

```
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan700 vif vlan700 address 7000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700
```

Step 3: Configure static route on switchA,B,C, So that the loopback interface could be reachable

switchA:

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 3000::2
```

switchB:

```
admin@XorPlus#set protocols static route 3131::31/128 next-hop 5000::1
admin@XorPlus#set protocols static route 2828::28/128 next-hop 3000::1
```

switchC:

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 5000::2
```

Step 4. Configure BGP

switchA:

```
admin@XorPlus#set protocols bgp bgp-id 28.28.28.28
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp fast-external-fallover disable true
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "2828::28"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp bgp-id 17.17.17.17
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus#set protocols bgp peer 2828::28 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 2828::28 as "17"
admin@XorPlus#set protocols bgp peer 2828::28 next-hop-self true
admin@XorPlus#set protocols bgp peer 2828::28 client true
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 3131::31 as "17"
admin@XorPlus#set protocols bgp peer 3131::31 next-hop-self true
admin@XorPlus#set protocols bgp peer 3131::31 client true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp bgp-id 31.31.31.31
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "3131::31"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

Step 5: Enable Labeled-BGP on switchA,B,C

switchA:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-label-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-label-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

Step 6: switchA distribute a bgp route

switchA:

```
admin@XorPlus#set protocols bgp network6 2001::/64
```

Step 7: Check BGP route table**switchA:**

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 :: nolabel/nolabel
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 2828::28 19/imp-null
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 1717::17 19/19
admin@XorPlus#
```

Step 8: Check mpls forwarding table on switchA,B,C**switchA:**

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 9: Configure MPLS label mapping on switchC

switchC:

```
admin@XorPlus#set protocols mpls in-label 1000 connect6 2001::/64
```

Note : The label 1000 is the top label that incoming BGP domain

Step 10:Configure MPLS label mapping on switchA

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop6 2001::2
```

Note : The label 2000 is the second-level label that incoming BGP domain , and it's the outgoing label that out of BGP domain the next-hop is indicate the next-hop ip address where the mpls packet forwarding and with a label 2000

Step 11:Check the mpls forwarding table on switchA ,B,C

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
2000 2000 -- -- 2001::2
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
1000 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 12:H2 send mpls packet with 3-level mpls label

Step13.Check the packet out of switchA, the top label 1000 will be popped at the Penultimate Hop

```

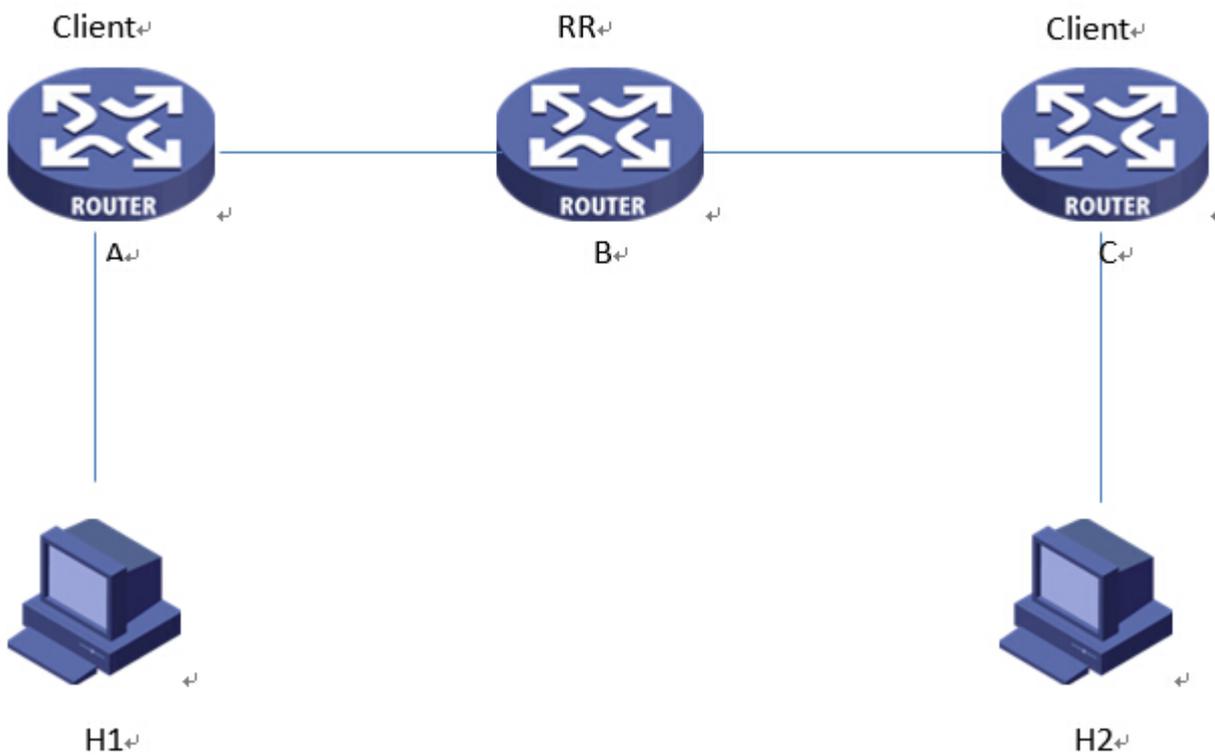
MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 00 03 00 00 00
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                = 0x8847 (Ethernet II)
MAC:

MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live       = 61 (0x3D)
MPLS: Label 2
MPLS: Label              = 2 (0x2)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live       = 64 (0x40)
MPLS:

IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class      = 003 (0x03)
IPv6: Flow Label         = 0000000 (0x000000)
IPv6: Payload Length     = 00000 (0x00000)
IPv6: Next Header        = 059 (0x3B)
IPv6: Hop Limit          = 255 (0xFF)
IPv6: Source Address:    7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:     2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Scenario 2: Incoming two layer label



Step 1: Configure Loopback interface

switchA:

```
admin@XorPlus#set vlan-interface loopback address 2828::28 prefix-length 128
```

switchB:

```
admin@XorPlus#set vlan-interface loopback address 1717::17 prefix-length 128
```

switchC:

```
admin@XorPlus#set vlan-interface loopback address 3131::31 prefix-length 128
```

Step 2:Configure Vlan-interface

switchA:

```
admin@XorPlus#set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlan-interface interface vlan200 vif vlan200 address 2001::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus#set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

switchB:

```
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::2
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::2
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300
```

switchC:

```
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan700 vif vlan700 address 7000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700
```

Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable

switchA:

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 3000::2
```

switchB:

```
admin@XorPlus#set protocols static route 3131::31/128 next-hop 5000::1
admin@XorPlus#set protocols static route 2828::28/128 next-hop 3000::1
```

switchC:

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 5000::2
```

Step 4.Configure BGP

switchA:

```
admin@XorPlus#set protocols bgp bgp-id 28.28.28.28
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp fast-external-fallover disable true
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "2828::28"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp bgp-id 17.17.17.17
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus#set protocols bgp peer 2828::28 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 2828::28 as "17"
admin@XorPlus#set protocols bgp peer 2828::28 next-hop-self true
admin@XorPlus#set protocols bgp peer 2828::28 client true
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 3131::31 as "17"
admin@XorPlus#set protocols bgp peer 3131::31 next-hop-self true
admin@XorPlus#set protocols bgp peer 3131::31 client true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp bgp-id 31.31.31.31
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "3131::31"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

Step 5:Enable Labeled-BGP on switchA,B,C:

switchA:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-label-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-label-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

Step 6:A distribute a bgp route

switchA:

```
admin@XorPlus#set protocols bgp network6 2001::/64
```

Step 7:Check BGP route table

switchA:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 :: nolabel/nolabel
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 2828::28 19/imp-null
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 1717::17 19/19
admin@XorPlus#
```

Step 8:Check mpls forwarding table on switchA,B,C:

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 9:Configure MPLS label mapping on switchC

switchC:

```
admin@XorPlus#set protocols mpls in-label 1000 connect6 2001::/64
```

Note : The label 1000 is the top label that incoming BGP domain

Step 10:Configure MPLS label mapping on switchA

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop6 2001::2
```

Note : The label 2000 is the second-level label that incoming BGP domain , and it's the outgoing label that out of BGP domain the next-hop is indicate the next-hop ip address where the mpls packet forwarding and with a label 2000

Step 11:Check the mpls forwarding table on A ,B,C

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
2000 2000 -- -- 2001::2
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
1000 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 12:H2 send mpls packet with 2-level mpls label

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 01 09 00 00 00
MAC: Type                = 0x8847 (Ethernet II)
MAC:

MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 1000 (0x3E8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 0 (0x0)
MPLS: Time To Live      = 64 (0x40)
MPLS: Label 2
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live      = 64 (0x40)
MPLS:

IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class     = 003 (0x03)
IPv6: Flow Label        = 0000000 (0x000000)
IPv6: Payload Length    = 00000 (0x00000)
IPv6: Next Header       = 059 (0x3B)
IPv6: Hop Limit         = 255 (0xFF)
IPv6: Source Address:   7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:    2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Step 13. Check the packet out of A, the top label 1000 will be popped at the Penultimate Hop

```

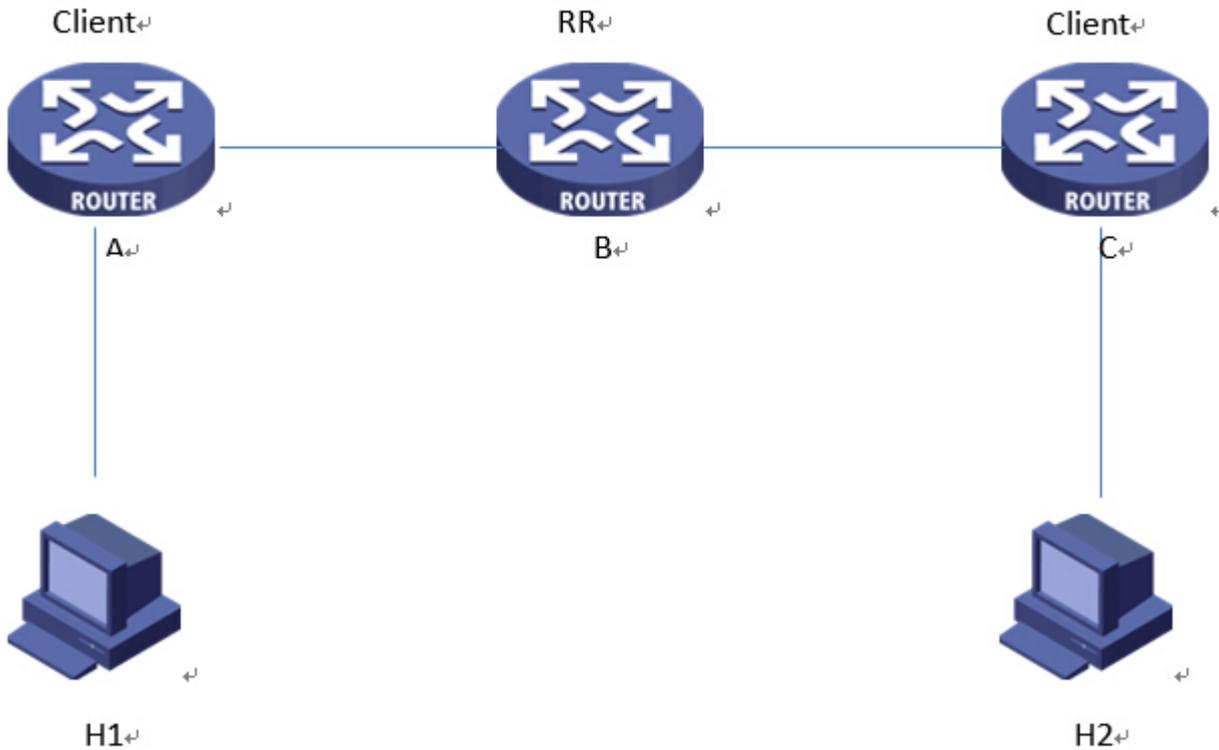
MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 00 03 00 00 00
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                = 0x8847 (Ethernet II)
MAC:

MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 2000 (0x7D0)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live      = 61 (0x3D)
MPLS:

IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class     = 003 (0x03)
IPv6: Flow Label        = 0000000 (0x000000)
IPv6: Payload Length    = 00000 (0x00000)
IPv6: Next Header       = 059 (0x3B)
IPv6: Hop Limit         = 255 (0xFF)
IPv6: Source Address:   7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:    2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Scenario 3: Incoming one layer label



Step 1: Configure Loopback interface

switchA:

```
admin@XorPlus#set vlan-interface loopback address 2828::28 prefix-length 128
```

switchB:

```
admin@XorPlus#set vlan-interface loopback address 1717::17 prefix-length 128
```

switchC:

```
admin@XorPlus#set vlan-interface loopback address 3131::31 prefix-length 128
```

Step 2: Configure Vlan-interface

switchA:

```
admin@XorPlus#set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlan-interface interface vlan200 vif vlan200 address 2001::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus#set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

switchB:

```

admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::2
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::2
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300

```

switchC:

```

admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan700 vif vlan700 address 7000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700

```

Step 3: Configure static route on switchA,B,C, So that the loopback interface could be reachable

switchA:

```

admin@XorPlus#set protocols static route 1717::17/128 next-hop 3000::2

```

switchB:

```

admin@XorPlus#set protocols static route 3131::31/128 next-hop 5000::1
admin@XorPlus#set protocols static route 2828::28/128 next-hop 3000::1

```

switchC:

```

admin@XorPlus#set protocols static route 1717::17/128 next-hop 5000::2

```

Step 4. Configure BGP

switchA:

```

admin@XorPlus#set protocols bgp bgp-id 28.28.28.28
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp fast-external-fallover disable true
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "2828::28"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true

```

switchB:

```

admin@XorPlus#set protocols bgp bgp-id 17.17.17.17
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus#set protocols bgp peer 2828::28 local-ip "1717::17"

```

```

admin@XorPlus#set protocols bgp peer 2828::28 as "17"
admin@XorPlus#set protocols bgp peer 2828::28 next-hop-self true
admin@XorPlus#set protocols bgp peer 2828::28 client true
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 3131::31 as "17"
admin@XorPlus#set protocols bgp peer 3131::31 next-hop-self true
admin@XorPlus#set protocols bgp peer 3131::31 client true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-unicast true

```

switchC:

```

admin@XorPlus#set protocols bgp bgp-id 31.31.31.31
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "3131::31"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true

```

Step 5: Enable Labeled-BGP on switchA,B,C:

switchA:

```

admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true

```

switchB:

```

admin@XorPlus#set protocols bgp peer 2828::28 ipv6-label-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-label-unicast true

```

switchC:

```

admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true

```

Step 6: A distribute a bgp route

switchA:

```

admin@XorPlus#set protocols bgp network6 2001::/64

```

Step 7: Check BGP route table

switchA:

```

admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 :: nolabel/nolabel
admin@XorPlus#

```

switchB:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 2828::28 19/imp-null
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 1717::17 19/19
admin@XorPlus#
```

Step 8: Check mpls forwarding table on switchA,B,C:

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 9: Configure MPLS label mapping on switchC

switchC:

```
admin@XorPlus#set protocols mpls in-label 1000 connect6 2001::/64
```

Note : The label 1000 is the top label that incoming BGP domain

Step 10: Configure MPLS label mapping on switchA

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop6 2001::2
```

Note : The label 2000 is the second-level label that incoming BGP domain , and it's the outgoing label that out of BGP domain the next-hop is indicate the next-hop ip address where the mpls packet forwarding and with a label 2000

Step 11: Check the mpls forwarding table on switchA,B,C

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
2000 2000 -- -- 2001::2
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
1000 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 12: H2 send mpls packet with 1-level mpls label

```
MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 01 09 00 00 00
MAC: Type                = 0x8847 (Ethernet II)
MAC:
MPLS: ----- MPLS Label -----
MPLS:
MPLS: Label 1
MPLS: Label              = 1000 (0x3E8)
MPLS: Experimental Use   = 0 (0x0)
MPLS: Bottom Of Stack (last entry) = 1 (0x1)
MPLS: Time To Live       = 64 (0x40)
MPLS:
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class      = 003 (0x03)
IPv6: Flow Label         = 0000000 (0x000000)
IPv6: Payload Length     = 00002 (0x00002)
IPv6: Next Header        = 059 (0x3B)
IPv6: Hop Limit          = 255 (0xFF)
IPv6: Source Address:    7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:     2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)
```

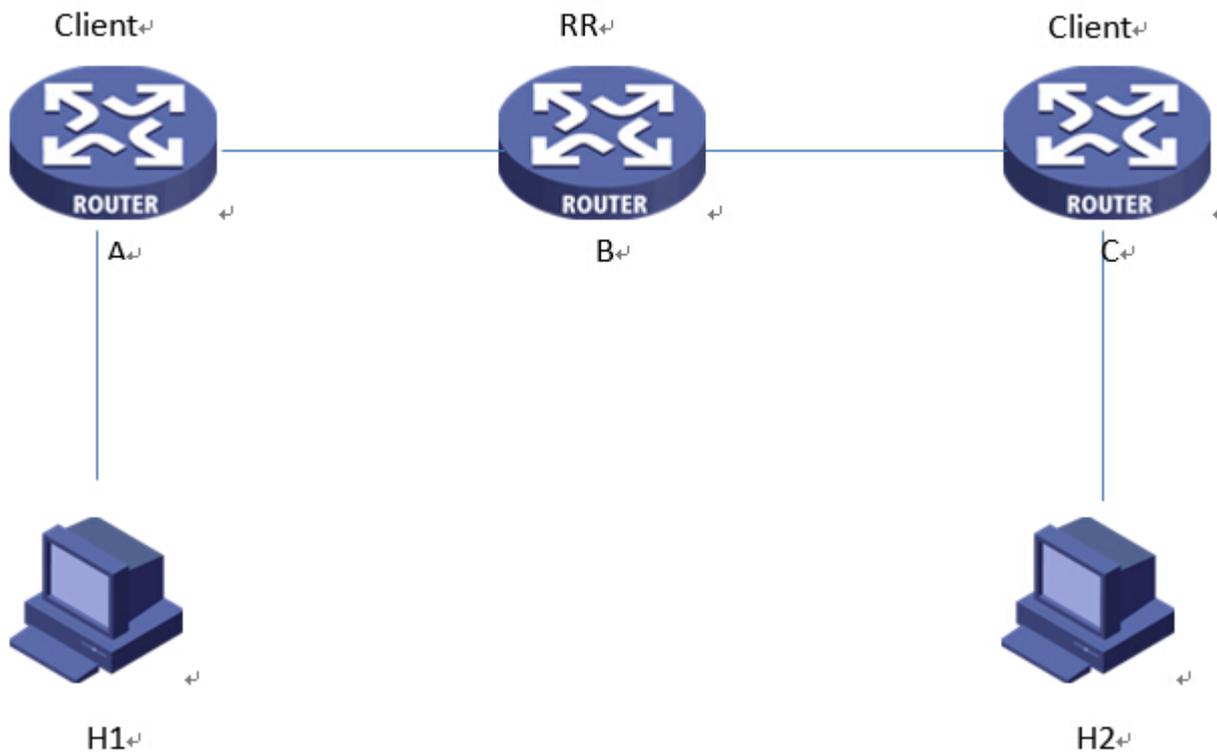
Step 13. Check the packet out of A, the top label 1000 will be popped at the Penultimate Hop

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 00 03 00 00 00
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                : 0x86DD (Ethernet II)
MAC:
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version              = 6 (0x6)
IPv6: Traffic Class        = 003 (0x03)
IPv6: Flow Label           = 0000000 (0x000000)
IPv6: Payload Length       = 00038 (0x0026)
IPv6: Next Header          = 059 (0x3B)
IPv6: Hop Limit            = 125 (0x7D)
IPv6: Source Address:      7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:        2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Scenario 4: Incoming IPv6 packet with no label



Step 1: Configure Loopback interface

switchA:

```
admin@XorPlus#set vlan-interface loopback address 2828::28 prefix-length 128
```

switchB:

```
admin@XorPlus#set vlan-interface loopback address 1717::17 prefix-length 128
```

switchC:

```
admin@XorPlus#set vlan-interface loopback address 3131::31 prefix-length 128
```

Step 2:Configure Vlan-interface**switchA:**

```
admin@XorPlus#set vlans vlan-id 200 l3-interface "vlan200"
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlan-interface interface vlan200 vif vlan200 address 2001::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/6 family ethernet-switching
native-vlan-id 300
admin@XorPlus#set interface gigabit-ethernet te-1/1/38 family ethernet-switching
native-vlan-id 200
```

switchB:

```
admin@XorPlus#set vlans vlan-id 300 l3-interface "vlan300"
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlan-interface interface vlan300 vif vlan300 address 3000::2
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::2
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/3 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/2 family ethernet-switching
native-vlan-id 300
```

switchC:

```
admin@XorPlus#set vlans vlan-id 500 l3-interface "vlan500"
admin@XorPlus#set vlans vlan-id 700 l3-interface "vlan700"
admin@XorPlus#set vlan-interface interface vlan500 vif vlan500 address 5000::1
prefix-length 64
admin@XorPlus#set vlan-interface interface vlan700 vif vlan700 address 7000::1
prefix-length 64
admin@XorPlus#set interface gigabit-ethernet te-1/1/46 family ethernet-switching
native-vlan-id 500
admin@XorPlus#set interface gigabit-ethernet te-1/1/48 family ethernet-switching
native-vlan-id 700
```

Step 3:Configure static route on switchA,B,C, So that the loopback interface could be reachable**switchA:**

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 3000::2
```

switchB:

```
admin@XorPlus#set protocols static route 3131::31/128 next-hop 5000::1
admin@XorPlus#set protocols static route 2828::28/128 next-hop 3000::1
```

switchC:

```
admin@XorPlus#set protocols static route 1717::17/128 next-hop 5000::2
```

Step 4.Configure BGP

switchA:

```
admin@XorPlus#set protocols bgp bgp-id 28.28.28.28
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp fast-external-fallover disable true
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "2828::28"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp bgp-id 17.17.17.17
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp route-reflector cluster-id 17.17.17.17
admin@XorPlus#set protocols bgp peer 2828::28 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 2828::28 as "17"
admin@XorPlus#set protocols bgp peer 2828::28 next-hop-self true
admin@XorPlus#set protocols bgp peer 2828::28 client true
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 local-ip "1717::17"
admin@XorPlus#set protocols bgp peer 3131::31 as "17"
admin@XorPlus#set protocols bgp peer 3131::31 next-hop-self true
admin@XorPlus#set protocols bgp peer 3131::31 client true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp bgp-id 31.31.31.31
admin@XorPlus#set protocols bgp local-as "17"
admin@XorPlus#set protocols bgp peer 1717::17 local-ip "3131::31"
admin@XorPlus#set protocols bgp peer 1717::17 as "17"
admin@XorPlus#set protocols bgp peer 1717::17 next-hop-self true
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-unicast true
```

Step 5:Enable Labeled-BGP on switchA,B,C:

switchA:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

switchB:

```
admin@XorPlus#set protocols bgp peer 2828::28 ipv6-label-unicast true
admin@XorPlus#set protocols bgp peer 3131::31 ipv6-label-unicast true
```

switchC:

```
admin@XorPlus#set protocols bgp peer 1717::17 ipv6-label-unicast true
```

Step 6:A distribute a bgp route

switchA:

```
admin@XorPlus#set protocols bgp network6 2001::/64
```

Step 7:Check BGP route table

switchA:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 :: nolabel/nolabel
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 2828::28 19/imp-null
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show bgp routes ipv6 label
Status Codes: * valid route, > best route
Prefix Nexthop In-label/Out-label
-----
*> 2001::/64 1717::17 19/19
admin@XorPlus#
```

Step 8:Check mpls forwarding table on switchA,B,C:

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 imp-null 2001::/64 te-1/1/2 3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label Out-label Destination Outgoing-Interface Next Hop
-----
19 19 2001::/64 te-1/1/44 5000::2
admin@XorPlus#
```

Step 9:Configure MPLS label mapping on switchC

switchC:

```
admin@XorPlus#set protocols mpls in-label 1000 connect6 2001::/64
```

Note : The label 1000 is the top label that incoming BGP domain

Step 10:Configure MPLS label mapping on switchA

```
admin@XorPlus#set protocols mpls in-label 2000 next-hop6 2001::2
```

Note : The label 2000 is the second-level label that incoming BGP domain , and it's the outgoing label that out of BGP domainthe next-hop is indicate the next-hop ip address where the mpls packet forwarding and with a label 2000.

Step 11:Check the mpls forwarding table on switchA ,B,C

switchA:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
2000      2000        --              --                  2001::2
admin@XorPlus#
```

switchB:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
19        imp-null   2001::/64       te-1/1/2            3000::1
admin@XorPlus#
```

switchC:

```
admin@XorPlus# run show mpls forwarding-table ipv6 all
In-label  Out-label  Destination      Outgoing-Interface  Next Hop
-----  -
19        19         2001::/64       te-1/1/44           5000::2
1000      19         2001::/64       te-1/1/44           5000::2
admin@XorPlus#
```

Step 12:H2 send IPv6 packet with no-label

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 04 7D 7B 62 93 FF
MAC: Source Address      : 00 01 09 00 00 00
MAC: Type                : 0x86DD (Ethernet II)
MAC:
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class     = 003 (0x03)
IPv6: Flow Label        = 0000000 (0x000000)
IPv6: Payload Length    = 00042 (0x002A)
IPv6: Next Header       = 059 (0x3B)
IPv6: Hop Limit         = 255 (0xFF)
IPv6: Source Address:   7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:    2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Step 13. Check the packet out of switchA

```

MAC: ----- MAC Header -----
MAC:
MAC: Destination Address : 00 00 03 00 00 00
MAC: Source Address      : 60 EB 69 9B BE 31
MAC: Type                : 0x86DD (Ethernet II)
MAC:
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version            = 6 (0x6)
IPv6: Traffic Class     = 003 (0x03)
IPv6: Flow Label        = 0000000 (0x000000)
IPv6: Payload Length    = 00042 (0x002A)
IPv6: Next Header       = 059 (0x3B)
IPv6: Hop Limit         = 252 (0xFC)
IPv6: Source Address:   7000:0000:0000:0000:0000:0000:0000:0002
IPv6: Dest Address:    2001:0000:0000:0000:0000:0000:0000:0002
IPv6:
IPv6: TransportLayer : Undecoded Protocol (59)

```

Debugging CLI for Labeled-BGP

Show BGP Peers

Note: This CLI is used to show the BGP peer status.

```

admin@XorPlus> show bgp peers
Peer 1: local 3131::31 remote 1717::17 state ESTABLISHED

```

Show BGP Peer Details

Note: This CLI is used to show BGP peer detail info.

```

admin@XorPlus> show bgp peers detail
Peer 1: local 3131::31 remote 1717::17 state ESTABLISHED
Peer ID: 17.17.17.17
Peer Type: IBGP
Admin State: START
BFD: Disabled
Address family: ipv4-unicast ipv6-unicast ipv6-label-unicast
Options: < Refresh >
Negotiated BGP Version: 4

```

```

Peer AS Number: 17
Updates Received: 5, Updates Sent: 0
Messages Received: 514, Messages Sent: 509
Time since last received update: 4994 seconds
Number of transitions to ESTABLISHED: 2
Time since last entering ESTABLISHED state: 13360 seconds
Retry Interval: 120 seconds
Hold Time: 90 seconds, Keep Alive Time: 30 seconds
Configured Hold Time: 90 seconds, Configured Keep Alive Time: 30 seconds
Minimum AS Origination Interval: 0 seconds
Minimum Route Advertisement Interval: 0 seconds
Received prefixes: 1
Accepted prefixes: 1
Active prefixes:

```

Show BGP Routes IPv4/IPv6

Note: This CLI is used to show BGP route table.

```

admin@XorPlus> show bgp routes ipv6
Status Codes: * valid route, > best route
Origin Codes: i IGP, e EGP, ? incomplete

```

Prefix	Nexthop	Peer	AS Path
*> 2001::/64	1717::17	17.17.17.17	i

Show BGP Routes IPv4/IPv6 Detail

Note: This CLI is used to show the detailed info of BGP route.

```

admin@XorPlus> show bgp routes ipv6 detail
2001::/64
  From peer: 17.17.17.17
  Route: Winner
  Origin: IGP
  AS Path:
  Nexthop: 1717::17
  Nexthop references: 1
  Local Preference: 100
  Originator ID: 28.28.28.28
  Cluster List: 17.17.17.17
  Age: 1:25:20

```

Show BGP Routes IPv4/IPv6 Label

Note: This CLI is used to show the label of the BGP speaker assigned to the specify BGP route entry.

```

admin@XorPlus> show bgp routes ipv6 label
Status Codes: * valid route, > best route

```

Prefix	Nexthop	In-label/Out-label
*> 2001::/64	1717::17	16/16

Show Route Label

Note: This CLI is used to show the label info of the route table.

```
admin@XorPlus> show route label
IPv4 Routing table: 4 routes
17.17.17.17/32    [static(1)/1]
> to 192.168.5.2 via vlan500/vlan500
  in-label:nolabel out-label:nolabel
31.31.31.31/32  [connected(0)/0]
> via loopback/loopback
  in-label:nolabel out-label:nolabel
192.168.5.0/30  [connected(0)/0]
> via vlan500/vlan500
  in-label:nolabel out-label:nolabel
192.168.7.0/30  [connected(0)/0]
> via vlan700/vlan700
  in-label:nolabel out-label:nolabel

IPv6 Routing table: 8 routes
1717::17/128    [static(1)/1]
> to 5000::2 via vlan500/vlan500
  in-label:nolabel out-label:nolabel
2001::/64       [ibgp(200)/0]
> to 5000::2 via vlan500/vlan500
  in-label:16 out-label:16
3131::31/128   [connected(0)/0]
> via loopback/loopback
  in-label:nolabel out-label:nolabel
5000::/64       [connected(0)/0]
> via vlan500/vlan500
  in-label:nolabel out-label:nolabel
7000::/64       [connected(0)/0]
> via vlan700/vlan700
  in-label:nolabel out-label:nolabel
fe80::/64       [connected(0)/0]
> via loopback/loopback
  in-label:nolabel out-label:nolabel
fe80::/64       [connected(0)/0]
> via vlan500/vlan500
  in-label:nolabel out-label:nolabel
fe80::/64       [connected(0)/0]
> via vlan700/vlan700
  in-label:nolabel out-label:nolabel
```

Show MPLS Forwarding-Table IPv4/IPv6 All

Note: This CLI is used to show the MPLS forwarding table.

```
admin@XorPlus> show mpls forwarding-table ipv6 all
In-label  Out-label  Destination          Outgoing-Interface  Next Hop
-----  -
16         16         2001::/64           te-1/1/44           5000::2
1000      16         2001::/64           te-1/1/44           5000::2
```

set protocols bgp trace options flag [all/

**FSM/configuration/damp/event/keepalive/nexthop-resolution/notification/notify-snmp/open
timer/update] disable false**

Note: This CLI is used to enable the debug switch for BGP.

set protocols mpls traceoptions disable false

Note: This CLI is used to enable mpls debug switch for mpls.

Layer 3 Multicast Routing Configuration

This chapter describes IGMP, PIM-SM, and IGMP snooping configurations.

- IGMP Snooping Configuration
- IGMP Configuration
- PIM-SM Configuration
- PIM-SM Configuration Example
- Multicast Command List
- PIM-SSM Configuration
- PIM-SSM Configuration Example

IGMP Snooping Configuration

In L2/L3, IGMPv2/IGMPv3 Snooping and IGMPv2 Snooping Querier are both supported.

IGMP Snooping Basic Configuration

In the default setting, the switch disables IGMP snooping. You should globally enable IGMP per VLAN.

```
admin@XorPlus# set protocols igmp-snooping enable true
admin@XorPlus# set protocols igmp-snooping vlan-id 1 enable true
admin@XorPlus# set protocols igmp-snooping vlan-id 1 mrouter interface ge-1/1/3
admin@XorPlus# set protocols igmp-snooping vlan-id 1 querier other-querier-timer 1
admin@XorPlus# set protocols igmp-snooping vlan-id 1 static group 238.255.0.1 interface
ge-1/1/2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show igmp-snooping vlan 1
Vlan 1:
-----
IGMP snooping : Enabled
IGMPv2 fast leave : Disabled
IGMP querier state : Disabled
IGMP querier source ip address : 0.0.0.0
IGMP other querier timer : 1
IGMP querier version : 2
admin@Xorplus# run show igmp-snooping groups
Vlan          Group          Port List          Type
-----
1             238.255.0.1    ge-1/1/2          Static
```

IGMP Snooping Querier

For multicast traffic in Layer2, enable an IGMP snooping querier in the VLAN.

```

admin@XorPlus# set protocols igmp-snooping vlan-id 1 querier enable true
admin@XorPlus# set protocols igmp-snooping vlan-id 1 querier address 10.10.1.1
admin@XorPlus# set protocols igmp-snooping vlan-id 1 querier version 2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show igmp-snooping querier
Vlan          IP Address    IGMP Version
-----
1             10.10.1.1    v2

```

IGMP Configuration

In admin@XorPlus, IGMPv1/v2/v3 is supported.

Configuring an IGMP Interface

Enable the multicast interface before enabling the IGMP interface.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan2 vif vlan2 address 10.10.60.10
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan3 vif vlan3 address 10.10.61.10
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan2 vif vlan2 disable false
admin@XorPlus# set multicast-interface interface vlan3 vif vlan3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols igmp interface vlan2 vif vlan2
admin@XorPlus# set protocols igmp interface vlan3 vif vlan3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show igmp interface
Interface State Querier Timeout Version Groups
-----
vlan2 UP 10.10.60.10 None 2 2
vlan3 UP 10.10.61.10 None 2 2
admin@XorPlus#

```

Configuring IGMP Parameters for the IGMP Interface

```
admin@XorPlus# set protocols igmp interface vlan2 vif vlan2 query-interval 4
admin@XorPlus# set protocols igmp interface vlan2 vif vlan2 query-last-member-interval 3
admin@XorPlus# set protocols igmp interface vlan2 vif vlan2 query-response-interval 100
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
```

Configuring an IGMPv3 Interface

You can configure IGMPv3 in a specified interface.

```
admin@XorPlus# set protocols igmp interface vlan3 vif vlan3 version 3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show igmp interface
Interface State Querier Timeout Version Groups
-----
vlan2 UP 10.10.60.10 None 2 2
vlan3 UP 10.10.61.10 None 3 2
```

Joining and Leaving a Group; Displaying Group Information

If you send an IGMPv2 report to VLAN 2 and an IGMPv3 report to VLAN 3, for example, you can display the group information of the switch. You should not have to worry about 224.0.0.2, 224.0.0.22, etc., which are used for the system (e.g. OSPF, RIP).

```
admin@XorPlus# run show igmp group
Interface Group Source LastReported Timeout V State
-----
vlan2 224.0.0.2 0.0.0.0 10.10.60.10 92 2 E
vlan2 224.0.0.22 0.0.0.0 10.10.60.10 101 2 E
vlan2 238.255.0.1 0.0.0.0 10.10.60.100 61 2 E
vlan3 224.0.0.2 0.0.0.0 10.10.61.10 205 3 E
vlan3 224.0.0.22 0.0.0.0 10.10.61.10 205 3 E
vlan3 238.255.0.2 0.0.0.0 10.10.61.100 0 3 I
vlan3 238.255.0.2 20.20.20.20 10.10.61.100 257 3 F
```

If you send a leaving message for the above group, the specified group will be removed.

```
admin@XorPlus# run show igmp group
Interface Group Source LastReported Timeout V State
-----
vlan2 224.0.0.2 0.0.0.0 10.10.60.10 88 2 E
vlan2 224.0.0.22 0.0.0.0 10.10.60.10 105 2 E
```

```
vlan3 224.0.0.2 0.0.0.0 10.10.61.10 227 3 E
vlan3 224.0.0.22 0.0.0.0 10.10.61.10 227 3 E
admin@XorPlus#
```

PIM-SM Configuration

In L2/L3, PIM-SM is supported.

PIM-SM Basic Configuration

Before configuring a PIM-SM interface, you should enable a multicast interface. You can then configure a candidate-RP and a candidate-BSR. For configuring the candidate-BSR, "scope-zone" denotes the zone of the multicast group, which is included in the multicast domain.

```
XorPlus# set vlans vlan-id 2 l3-interface vlan-2
XorPlus# set vlans vlan-id 3 l3-interface vlan-3
XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 2
XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 3
XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.60.10 prefix-length
24
XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.61.10 prefix-length
24
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set protocols igmp interface vlan-3 vif vlan-3
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set protocols pimsm4 bootstrap cand-bsr scope-zone 224.0.0.0/4
cand-bsr-by-vif-name vlan-3
XorPlus# set protocols pimsm4 bootstrap cand-rp group-prefix 237.0.0.0/8
cand-rp-by-vif-name vlan-2
XorPlus# set protocols pimsm4 bootstrap cand-rp group-prefix 231.0.0.0/8
cand-rp-by-vif-name vlan-3
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus#
```

Static RP Configuration

You can also configure static RP instead of BSR or dynamic RP.

```
XorPlus# set protocols pimsm4 static-rps rp 10.10.60.10 group-prefix 238.0.0.0/8
rp-priority 10
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus#
```

PIM-SM Configuration Example

In the following topology, Switch B is the C-BSR and C-RP. Host A is a receiver for multicast traffic, and Host B is a multicast source that will send the multicast traffic.

You'll need to configure ge-1/1/2 as an IGMP interface in Switch A for Host A.

In this example, the static route in the RIB will be used by PIM-SM.

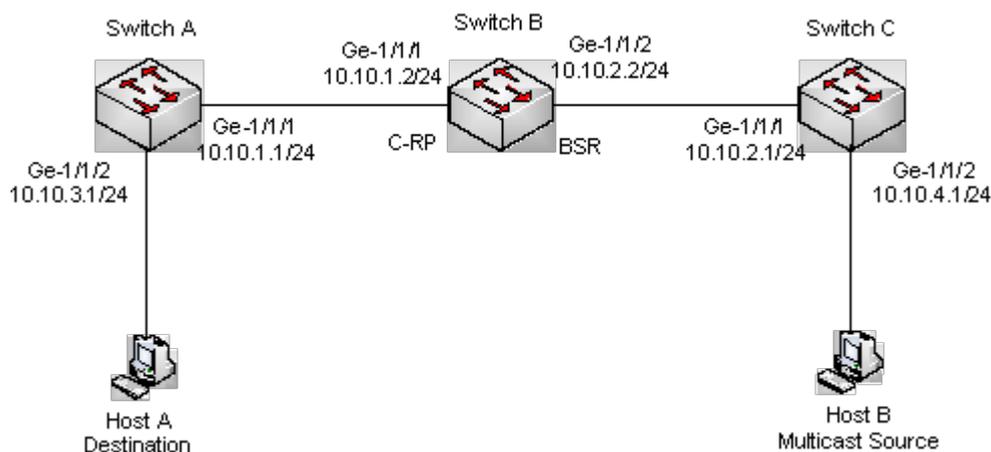


Figure 8-1. PIM-SM multicast routing configuration.

Configuring Switch A

For switch A, configure ge-1/1/2 as an IGMP interface and ge-1/1/1 as a PIM-SM interface.

```
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.3.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
```

```

admin@XorPlus# set protocols igmp interface vlan-3 vif vlan-3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#set protocols static route 10.10.2.0/24 next-hop 10.10.1.2
admin@XorPlus#set protocols static route 10.10.4.0/24 next-hop 10.10.1.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show pim interface
Interface State Mode V PIMstate Priority DRaddr Neighbors
-----
vlan-2 UP Sparse 2 DR 1 10.10.1.1 0
vlan-3 UP Sparse 2 DR 1 10.10.3.1 0
register_vif UP Sparse 2 DR 1 10.10.1.1 0
admin@XorPlus#
admin@XorPlus# run show igmp interface
Interface State Querier Timeout Version Groups
-----
vlan-2 DISABLED 10.10.1.1 None 2 0
vlan-3 UP 10.10.3.1 None 2 3

```

Configuring Switch B

Configure 2 PIM-SM interfaces, ge-1/1/1 and ge-1/1/2. You will also need to configure a candidate BSR and a candidate RP.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.2.2
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 bootstrap cand-bsr scope-zone 224.0.0.0/4
cand-bsr-by-vif-name vlan-3

```

```

admin@XorPlus# set protocols pimsm4 bootstrap cand-rp group-prefix 238.0.0.0/8
cand-rp-by-vif-name vlan-2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#set protocols static route 10.10.3.0/24 next-hop 10.10.1.1
admin@XorPlus#set protocols static route 10.10.4.0/24 next-hop 10.10.2.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show pim bootstrap
Active zones:
BSR Pri LocalAddress Pri State Timeout SZTimeout
10.10.2.2 1 10.10.2.2 1 Elected 19 -1
Expiring zones:
BSR Pri LocalAddress Pri State Timeout SZTimeout
admin@XorPlus#
admin@XorPlus# run show pim rps
RP Type Pri Holdtime Timeout ActiveGroups GroupPrefix
-----
10.10.1.2 bootstrap 192 150 -1 0 238.0.0.0/8
admin@XorPlus#

```

Configuring Switch C

Configure 2 PIM-SM interfaces, ge-1/1/1 and ge-1/1/2. You will also need to configure a candidate BSR and a candidate RP.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2vif vlan-2 address 10.10.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3vif vlan-3 address 10.10.4.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 10.10.1.0/24 next-hop 10.10.2.2
admin@XorPlus# set protocols static route 10.10.3.0/24 next-hop 10.10.2.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#

```

```

admin@XorPlus# run show pim interface
Interface State Mode V PIMstate Priority DRaddr Neighbors
-----
vlan-2 UP Sparse 2 DR 1 10.10.2.1 0
vlan-3 UP Sparse 2 DR 1 10.10.4.1 0
register_vif UP Sparse 2 DR 1 10.10.2.1 0
admin@XorPlus#

```

Multicast Command List

View the list of multicast commands supported by the Pica8 PicOS software here.

PIM-SSM Configuration

In L2/L3, PIM-SM is supported.

PIM-SM Basic Configuration

Before configuring a PIM-SM interface, user should enable a multicast interface.

```

XorPlus# set vlans vlan-id 2 l3-interface vlan-2
XorPlus# set vlans vlan-id 3 l3-interface vlan-3
XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 2
XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 3
XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.60.10 prefix-length
24
XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.61.10 prefix-length
24
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set protocols igmp interface vlan-3 vif vlan-3
XorPlus# set protocols igmp interface vlan-3 vif vlan-3 version 3
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

PIM-SSM Configuration Example

In the following topology, Host B is a multicast source that will send the multicast traffic.

You'll need to configure ge-1/1/2 as an IGMP interface in Switch A for Host A.

In this example, the static route in the RIB will be used by PIM-SSM.

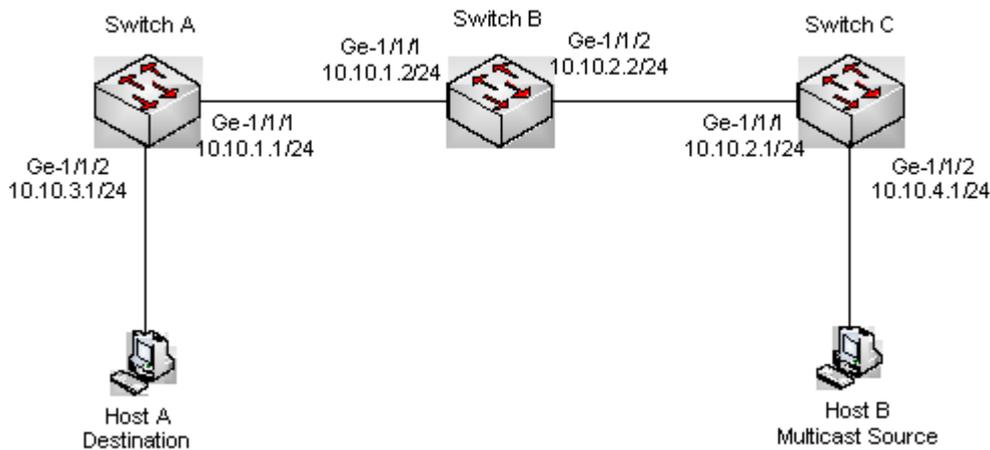


Figure 8-1. PIM-SSM multicast routing configuration.

Configuring Switch A

For Switch A, configure ge-1/1/2 as an IGMP interface, and ge-1/1/1 as a PIM-SM interface.

```
admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.3.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
admin@XorPlus# set protocols igmp interface vlan-3 vif vlan-3
admin@XorPlus# set protocols igmp interface vlan-3 vif vlan-3 version 3
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 10.10.2.0/24 next-hop 10.10.1.2
admin@XorPlus# set protocols static route 10.10.4.0/24 next-hop 10.10.1.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# run show pim interface
```

```

Interface State Mode V PIMstate Priority DRaddr Neighbors
-----
vlan-2 UP Sparse 2 DR 1 10.10.1.1 0
vlan-3 UP Sparse 2 DR 1 10.10.3.1 0
register_vif UP Sparse 2 DR 1 10.10.1.1 0
admin@XorPlus#
admin@XorPlus# run show igmp interface
Interface State Querier Timeout Version Groups
-----
vlan-2 DISABLED 10.10.1.1 None 2 0
vlan-3 UP 10.10.3.1 None 2 3

```

Configuring Switch B

Configure 2 PIM-SSM interfaces, ge-1/1/1 and ge-1/1/2.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.1.2
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.2.2
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#set protocols static route 10.10.3.0/24 next-hop 10.10.1.1
admin@XorPlus#set protocols static route 10.10.4.0/24 next-hop 10.10.2.1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.

```

Configuring Switch C

Configure 2 PIM-SSM interfaces, ge-1/1/1 and ge-1/1/2.

```

admin@XorPlus# set vlans vlan-id 2 l3-interface vlan-2
admin@XorPlus# set vlans vlan-id 3 l3-interface vlan-3
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching
native-vlan-id 2

```

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching
native-vlan-id 3
admin@XorPlus# set vlan-interface interface vlan-2 vif vlan-2 address 10.10.2.1
prefix-length 24
admin@XorPlus# set vlan-interface interface vlan-3 vif vlan-3 address 10.10.4.1
prefix-length 24
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set multicast-interface interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set multicast-interface interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols pimsm4 interface vlan-2 vif vlan-2 disable false
admin@XorPlus# set protocols pimsm4 interface vlan-3 vif vlan-3 disable false
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus# set protocols static route 10.10.1.0/24 next-hop 10.10.2.2
admin@XorPlus# set protocols static route 10.10.3.0/24 next-hop 10.10.2.2
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# run show pim interface
Interface State Mode V PIMstate Priority DRaddr Neighbors
-----
vlan-2 UP Sparse 2 DR 1 10.10.2.1 0
vlan-3 UP Sparse 2 DR 1 10.10.4.1 0
register_vif UP Sparse 2 DR 1 10.10.2.1 0
admin@XorPlus#
```

QoS Configuration

This chapter describes Layer2 and Layer3 QoS configurations.

- Weighted Random Early Detection
 - WRED Overview
 - WRED Configuration Tasks
 - WRED Configuration Example
- QoS Principle
- SP Configuration Example
- WRR Configuration Example
- WFQ Configuration Example
- QoS Configuration Guide
- PFC Configuration Example

Weighted Random Early Detection

WRED (Weighted Random Early Detection) is a congestion avoidance mechanism that makes use of the congestion control mechanism of TCP (Transmission Control Protocol). By selectively dropping packets before periods of congestion, WRED tells TCP senders to reduce their transmission rates.

This chapter provides an overview of WRED, describes WRED configuration tasks, and presents a configuration example.

- WRED Overview
- WRED Configuration Tasks
- WRED Configuration Example

WRED Overview

You can monitor the network traffic load on an interface to anticipate congestion. Congestion can be avoided by dropping packets when needed. Tail drop is the most basic mechanism for avoiding congestion. Interface output queues fill during periods of congestion. When an output queue is full, the tail drop mechanism kicks in, and packets are dropped from the tail of the queue until congestion is reduced and the queue is no longer full. Tail drop does not differentiate between classes of traffic and drops packets at the tail of an output queue indiscriminately.

Global synchronization is a problem that occurs when tail drop is used as the congestion avoidance mechanism. Tail drop results in a large number of packets getting dropped at once. In response to packet drops, multiple senders of TCP (Transmission Control Protocol) traffic reduce their transmission rates at around the same time. The same TCP senders then increase their transmission rates around the same time when congestion is reduced. It results in periods of high link utilization followed by periods of low utilization.

WRED (Weighted Random Early Detection) is a congestion avoidance mechanism that prevents the problem of global synchronization associated with tail drop. When an output queue begins to experience congestion, WRED starts dropping packets *selectively*. A TCP sender experiencing packet drops reduces its transmission rate. By dropping some packets earlier than the point when the queue is full, WRED prevents the situation where a large number of packets get dropped at once. WRED not only reduces the chances of global synchronization, it also increases the utilization of transmission bandwidth.

WRED should be configured on egress ports.

WRED Configuration Tasks

When a packet arrives at a WRED-enabled output interface, the following chain of events takes place:

1. The length of the queue is calculated.
2. If the queue length is less than the minimum threshold, the packet is placed in the queue.
3. If the queue length is more than the maximum threshold, the packet is dropped.
4. If the queue length is more than the minimum threshold but less than the maximum threshold, the packet is either dropped or queued, based on the packet drop probability.

The following command can be used to enable or disable WRED:

```
set interface gigabit-ethernet <port> wred queue <value> enable <bool>
```

The following example demonstrates how to enable WRED on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 enable true
admin@XorPlus# commit
Commit OK.
Save done.
```

The following example demonstrates how to disable WRED on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 enable false
admin@XorPlus# commit
Commit OK.
Save done.
```

The following command can be used to set the maximum threshold:

```
set interface gigabit-ethernet <port> wred queue <value> max_thresh <int>
```

The following example demonstrates how to set the maximum threshold to 400 on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 max_thresh 400
admin@XorPlus# commit
Commit OK.
Save done.
```

The following command can be used to set the minimum threshold:

set interface gigabit-ethernet <port> wred queue <value> min_thresh <int>

The following example demonstrates how to set the minimum threshold to 200 on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 min_thresh 200
admin@XorPlus# commit
Commit OK.
Save done.
```

The following command can be used to to configure drop probability:

set interface gigabit-ethernet <port> wred queue <value> drop_probability <int>

The following example demonstrates how to set the drop probability to 50% on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 drop_probability 50
admin@XorPlus# commit
Commit OK.
Save done.
```

When congestion occurs, WRED drops packets based on the queue length exceeding certain threshold value. ECN (Explicit Congestion Notification) can enhance basic WRED operation by marking packets instead of dropping them when the queue length exceeds certain threshold value. Downstream routers and hosts would see this marking as an indication of network congestion and slow down their packet transmission rates.

ECN is a value in the DS (Differentiated Services) field of the IPv4 protocol header. ECN uses the two least significant (right-most) bits of the 8-bit DF field to encode four different codepoints:

1. 00 - Not ECN-Capable Transport
2. 01 - ECN-Capable Transport(1)
3. 10 - ECN-Capable Transport(0)
4. 11 - Congestion Experienced

When both end hosts support ECN, they mark their packets with either 10 or 01. When ECN is enabled, PicOS changes the ECN field of all such packets to 11. When ECN is not enabled, the ECN bits are not changed.

The following command can be used to enable or disable ECN:

set interface gigabit-ethernet <port> wred queue <value> ecn_thresh <int>

The following example demonstrates how to enable ECN on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 ecn_thresh 1
admin@XorPlus# commit
Commit OK.
Save done.
```

The following example demonstrates how to disable ECN on queue 0 of interface ge-1/1/1:

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 wred queue 0 ecn_thresh 0
admin@XorPlus# commit
Commit OK.
Save done.
```

WRED Configuration Example

As shown in Fig 1, interfaces ge-1/1/1, ge-1/1/2, and ge-1/1/3 are connected to end hosts PC1, PC2, and PC3, respectively. Ports ge-1/1/1 and ge-1/1/2 are ingress ports, while ge-1/1/3 is an egress port. In this example, WRED is configured on ge-1/1/3.

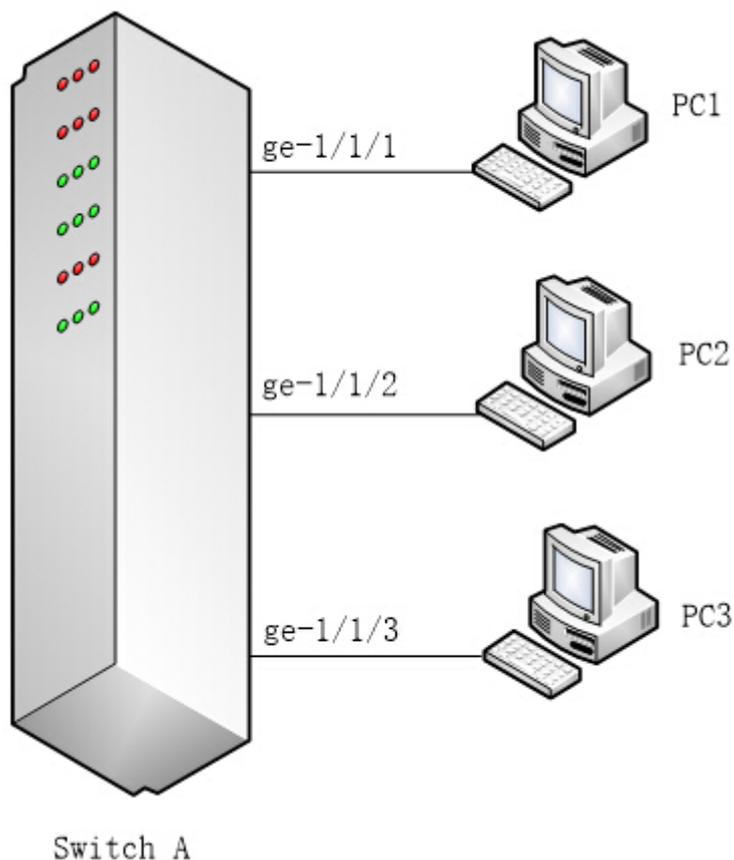


Fig 1. Configure WRED

Enable WRED

Enable WRED on queue 0 of interface ge-1/1/3.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 wred queue 0 enable true
admin@XorPlus# commit
Commit OK.
Save done.
```

Set Maximum and Minimum Thresholds

Set the maximum threshold to 400 and the minimum threshold to 200 on queue 0 of interface ge-1/1/3.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 wred queue 0 max_thresh 400
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 wred queue 0 min_thresh 200
admin@XorPlus# commit
Commit OK.
Save done.
```

Set Drop Probability

Set the drop probability to 50% on queue 0 of interface ge-1/1/3.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 wred queue 0 drop_probability 50
admin@XorPlus# commit
Commit OK.
Save done.
```

Enable ECN

Enable ECN (Explicit Congestion Notification) on queue 0 of interface ge-1/1/3.

```
admin@XorPlus# set interface gigabit-ethernet ge-1/1/3 wred queue 0 ecn_thresh 1
admin@XorPlus# commit
Commit OK.
Save done.
```

SP Queue Principle

When the scheduler mode is SP, the egress port has eight queues: 7, 6, 5, 4, 3, 2, 1, and 0, queue 7 being the highest priority and queue 0 being the lowest priority. The advantage is that it can give priority to the transmissions of a key business group. This scheduler mode also has disadvantages. During times of congested conditions, if higher the priority queue has groups for a long time, the low priority queues will not get service all the time.

WRR Principle

The full name of WRR is Weighted Round Robin. In order to ensure that every queue has certain servicing time, WRR uses a round robin scheduling algorithm between the queues. When the scheduler mode is WRR, every queue can have a weighted value, which is also known as scheduling weight. Scheduling weight means that when the egress port schedules, the queue messages the proportion of scheduling resources to be used. Scheduling unit is Kbps. The example of a WRR scheduling algorithm is as follows:

On the 1000 Mbps egress port, the scheduling weights of eight queues are 5, 4, 3, 3, 2, 1, 1, 1; this ensures that even the lowest priority queue gets bandwidth.

The calculation method is as follows:

$1/(5+4+3+3+2+1+1+1)*1000 \text{ Mbps}=50 \text{ Mbps}$.

This can avoid the problem of the packets in the lower priority queues not getting service for a long time. The advantage is that although the queue scheduling uses round robin scheduling, every queue does not distribute at a fixed service time—if a queue is empty, the next queue should be scheduled immediately. In this way, it makes full use of bandwidth resources. When using WRR scheduling mode, user can define the weighted value for each queue.

WFQ Principle

The full name of WFQ is Weighted Fair Queuing. It is similar to WRR. The only difference between WFQ and WRR is that the scheduling mode in WFQ supports a minimum bandwidth guarantee, making this scheduling scheme more flexible. Configuring a minimum guaranteed bandwidth assures that every queue working in WFQ mode has a minimum bandwidth guarantee. In addition, the bandwidth available for distribution allocates according to the weighted proportion in the corresponding queue.

The distributable bandwidth calculation method is as follows:

distributable bandwidth = total bandwidth - minimum bandwidth

The example of the WFQ scheduling algorithm as follows:

Assuming that the total bandwidth of the egress port is 100M, there are 3 flows in the queue of this port.

Their scheduling weighted values are 1,2,4; the minimum bandwidth guarantees of these 3 flows are 10000Kbps,10000Kbps, and 20000Kbps.

Proportions of each flow are 10%,10%, 20%.

Distributable bandwidth = 100M-(10M+10M+20M)=60M.

Proportion of distributable bandwidth is 60%

Total distributable bandwidth = the sum of each flow weighted value.

In this example, the total distributable bandwidth is 7(that is 1+2+4).

Formula to calculate the proportion of distributable bandwidth which is occupied by each flow is as follows:

Proportion of distributable bandwidth = (the own weight of flow)/(distributable bandwidth).

Proportions of the distributable bandwidth for each flow are 1/7, 2/7, 4/7.

Bandwidth ratio of the flows is (10%+60%*(1/7)) : (10%+60%*(2/7)) : (20%+60%*(4/7)),that is 13:19:38.

SP Configuration Example

As shown in Fig 1, ge-1/1/1 and ge-1/1/2 are ingress ports, and ge-1/1/3 is an egress port. Use default scheduling model: priority trust model is IEEE 802.1.

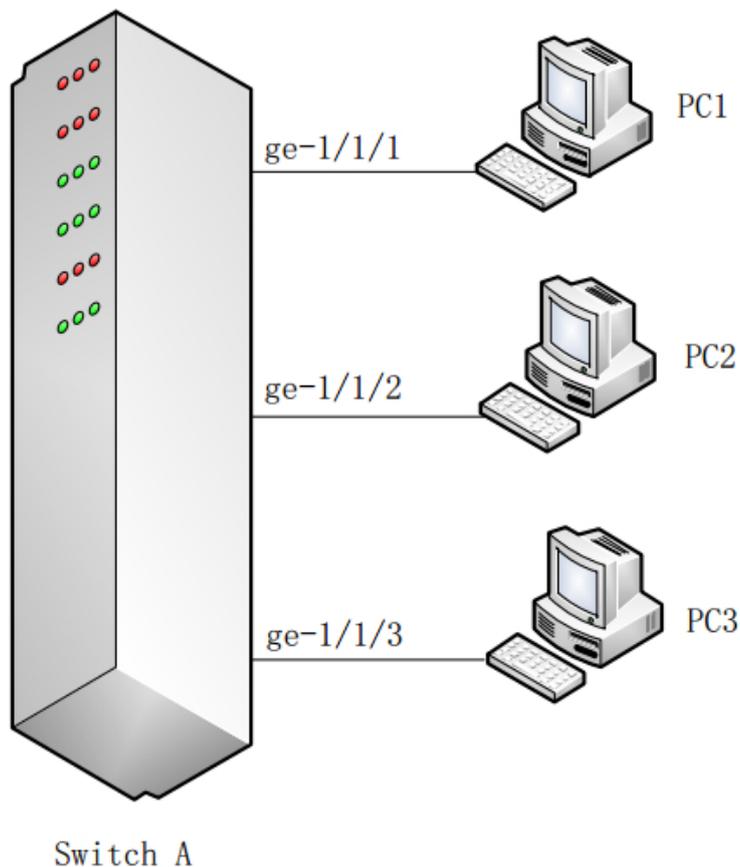


Fig 1. Configure SP

Configure two forwarding-classes

Configure forwarding-class f1 and f2 and their local-priorities.

```
admin@XorPlus# set class-of-service forwarding-class f1 local-priority 3
admin@XorPlus# set class-of-service forwarding-class f2 local-priority 6
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring classifier

Configure classifier c1, c2, and trust mode. Also configure classifier relevant to forwarding class and code point.

```
admin@XorPlus# set class-of-service classifier c1 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c1 forwarding-class f1 code-point 5
admin@XorPlus# set class-of-service classifier c2 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c2 forwarding-class f2 code-point 7
admin@XorPlus# commit
Commit OK.
Save done.
```

Apply classifiers to two ingress ports

Configure classifier c1 and apply it to port ge-1/1/1. Configure classifier c2 and apply it to port ge-1/1/2.

```
admin@XorPlus# set class-of-service interface ge-1/1/1 classifier c1
admin@XorPlus# set class-of-service interface ge-1/1/2 classifier c2
admin@XorPlus# commit
```

Commit OK.

Save done.

```
admin@XorPlus# run show class-of-service interface ge-1/1/1
Interface : ge-1/1/1
```

trust mode : ieee-802.1

Default ieee-802.1 : 0

Default dscp : 0

Default inet-precedence : 0

Local-priority	Queue-Schedule	Code-points
0	SP,0kbps	0
1	SP,0kbps	1
2	SP,0kbps	2
3	SP,0kbps	3,5
4	SP,0kbps	4
5	SP,0kbps	
6	SP,0kbps	6
7	SP,0kbps	7

```
admin@XorPlus# run show class-of-service interface ge-1/1/2
Interface : ge-1/1/2
```

trust mode : ieee-802.1

Default ieee-802.1 : 0

Default dscp : 0

Default inet-precedence : 0

Local-priority	Queue-Schedule	Code-points
0	SP,0kbps	0
1	SP,0kbps	1
2	SP,0kbps	2
3	SP,0kbps	3
4	SP,0kbps	4
5	SP,0kbps	5
6	SP,0kbps	6,7
7	SP,0kbps	

Generate Traffic

PC1 and PC2 generates traffic, which is matched with the corresponding classifier. Port PC1 and PC2 send 100% traffic to PC3 at the same time.

The expected result is that PC3 only can receive packets from PC2.

WRR Configuration Example

As shown in Fig 2, ge-1/1/1 and ge-1/1/2 are ingress ports. ge-1/1/3 is the egress port. Use WRR scheduling model: priority trust mode is IEEE 802.1.

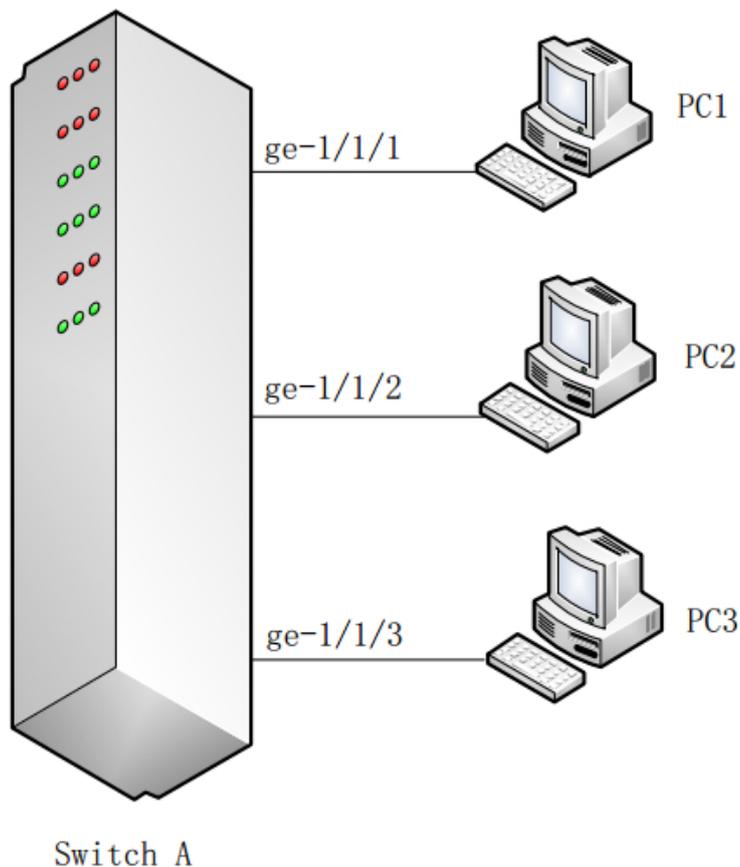


Fig 2. Configure WRR

Configure Scheduler

Configure two schedulers: s1 and s2. Their modes are WRR. Configure scheduler s1's weight as 1 and s2's weight as 3.

```
admin@XorPlus# set class-of-service scheduler s1 mode WRR
admin@XorPlus# set class-of-service scheduler s2 mode WRR
admin@XorPlus# set class-of-service scheduler s1 weight 1
admin@XorPlus# set class-of-service scheduler s2 weight 3
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure Two Forwarding-Classes

Configure forwarding-classes, f1, f2, and their local-priorities.

```
admin@XorPlus# set class-of-service forwarding-class f1 local-priority 3
admin@XorPlus# set class-of-service forwarding-class f2 local-priority 6
admin@XorPlus# commit
Commit OK.
Save done.
```

Configuring Classifiers and Apply Classifiers to Ingress Ports

Configure classifier c1, c2, and the trust mode for each. Configure classifier relevant to the forwarding class. Both c1 and c2 are used as ingress ports, and they should contain code point, not scheduler.

```

admin@XorPlus# set class-of-service classifier c1 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c1 forwarding-class f1 code-point 5
admin@XorPlus# set class-of-service classifier c2 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c2 forwarding-class f2 code-point 7
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set class-of-service interface ge-1/1/1 classifier c1
admin@XorPlus# set class-of-service interface ge-1/1/2 classifier c2
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show class-of-service interface ge-1/1/1
Interface : ge-1/1/1
trust mode : ieee-802.1
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule           Code-points
-----
0                SP,0kbps                 0
1                SP,0kbps                 1
2                SP,0kbps                 2
3                SP,0kbps                 3,5
4                SP,0kbps                 4
5                SP,0kbps
6                SP,0kbps                 6
7                SP,0kbps                 7
admin@XorPlus# run show class-of-service interface ge-1/1/2
Interface : ge-1/1/2
trust mode : ieee-802.1
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule           Code-points
-----
0                SP,0kbps                 0
1                SP,0kbps                 1
2                SP,0kbps                 2
3                SP,0kbps                 3
4                SP,0kbps                 4
5                SP,0kbps                 5
6                SP,0kbps                 6,7
7                SP,0kbps

```

Configuring Scheduler Profile and Apply Scheduler Profile to Egress Ports

Scheduler profile p1 is used for egress port ge-1/1/3 and should contain a scheduler not containing code point.

```

admin@XorPlus# set class-of-service scheduler-profile p1 forwarding-class f1 scheduler s1
admin@XorPlus# set class-of-service scheduler-profile p1 forwarding-class f2 scheduler s2
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set class-of-service interface ge-1/1/3 scheduler-profile p1

```

```

admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show class-of-service interface ge-1/1/3
Interface : ge-1/1/3
trust mode : no-trust
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule           Code-points
-----
0                SP,0kbps
1                SP,0kbps
2                SP,0kbps
3                WRR,1,0kbps
4                SP,0kbps
5                SP,0kbps
6                WRR,3,0kbps
7                SP,0kbps

```

Generate Traffic

PC1 and PC2 generate traffic that are matched with the corresponding classifier. PC1 and PC2 send 100% of traffic to PC3 at the same time.

The expected result is that PC3 can receive packets from PC1 and PC2, and their rate is 1:3, that is the weight proportion in the corresponding queue.

WFQ Configuration Example

As shown in Fig 3, ge-1/1/1 and ge-1/1/2 are ingress ports. ge-1/1/3 is the egress port. Use the WFQ scheduling model. Priority trust model is IEEE 802.1. The bandwidth is 100Mbps.

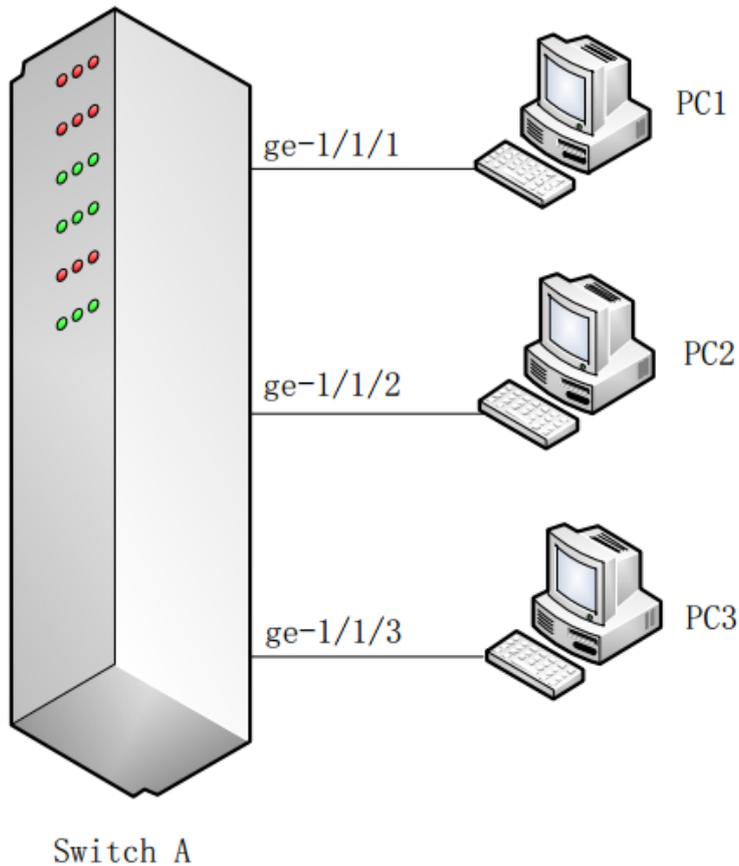


Fig 3. Configure WFQ

Configure Scheduler

Configure two schedulers (s1 and s2) with guaranteed-rates and their modes WFQ. Configure scheduler s1 weighted at 1 and scheduler s2 weighted at 3. The guaranteed-rate of scheduler s1 is 10000, and the guaranteed-rate of scheduler s2 is 30000.

```
admin@XorPlus# set class-of-service scheduler s1 mode WFQ
admin@XorPlus# set class-of-service scheduler s2 mode WFQ
admin@XorPlus# set class-of-service scheduler s1 weight 1
admin@XorPlus# set class-of-service scheduler s1 guaranteed-rate 10000
admin@XorPlus# set class-of-service scheduler s2 weight 3
admin@XorPlus# set class-of-service scheduler s2 guaranteed-rate 30000
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure Two Forwarding-Classes

Configure forwarding-class f1 and f2 and their local-priorities.

```

admin@XorPlus# set class-of-service forwarding-class f1 local-priority 3
admin@XorPlus# set class-of-service forwarding-class f2 local-priority 6
admin@XorPlus# commit
Commit OK.
Save done.

```

Configuring Classifier and Apply Classifiers to Ingress Ports

Configure classifier c1, c2, and c3, as well as the trust mode. Configure classifiers relevant to the forwarding class. Classifiers c1 and c2 are used as the ingress ports, and they should contain code point, not scheduler.

```

admin@XorPlus# set class-of-service classifier c1 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c1 forwarding-class f1 code-point 5
admin@XorPlus# set class-of-service classifier c2 trust-mode ieee-802.1
admin@XorPlus# set class-of-service classifier c2 forwarding-class f2 code-point 7
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set class-of-service interface ge-1/1/1 classifier c1
admin@XorPlus# set class-of-service interface ge-1/1/2 classifier c2
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show class-of-service interface ge-1/1/1
Interface : ge-1/1/1

trust mode : ieee-802.1
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule           Code-points
-----
0                SP,0kbps                 0
1                SP,0kbps                 1
2                SP,0kbps                 2
3                SP,0kbps                 3,5
4                SP,0kbps                 4
5                SP,0kbps
6                SP,0kbps                 6
7                SP,0kbps                 7
admin@XorPlus# run show class-of-service interface ge-1/1/2
Interface : ge-1/1/2

trust mode : ieee-802.1
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule           Code-points
-----
0                SP,0kbps                 0
1                SP,0kbps                 1
2                SP,0kbps                 2
3                SP,0kbps                 3
4                SP,0kbps                 4
5                SP,0kbps                 5
6                SP,0kbps                 6,7
7                SP,0kbps

```

Configuring Scheduler Profile and Apply Classifiers to Egress Ports

Scheduler profile p1 is used to egress port ge -1/1/3 and should contain a scheduler not containing code point.

```

admin@XorPlus# set class-of-service scheduler-profile p1 forwarding-class f1 scheduler s1
admin@XorPlus# set class-of-service scheduler-profile p1 forwarding-class f2 scheduler s2
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# set class-of-service interface ge-1/1/3 scheduler-profile p1
admin@XorPlus# commit
Commit OK.
Save done.

admin@XorPlus# run show class-of-service interface ge-1/1/3
Interface : ge-1/1/3

trust mode : no-trust
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0
Local-priority   Queue-Schedule                               Code-points
-----
0                SP,0kbps
1                SP,0kbps
2                SP,0kbps
3                WFQ,1,10000kbps
4                SP,0kbps
5                SP,0kbps
6                WFQ,3,30000kbps
7                SP,0kbps

```

Generate Traffic

PC1 and PC2 generate traffic which are matched with the corresponding classifier. PC1 and PC2 send 100% traffic to PC3 at the same time.

The expected result is that PC3 can receive packets from PC1 and PC2, and their rate is about 1:3. That is, the weight proportion and the guaranteed-rate have corresponding queues.

QoS Configuration Guide

Configure a Scheduler

A scheduler, which determines the QoS working mode and weight, should be configured first when you configure QoS. The working mode can be SP, WRR, or WFQ, and the weight is 1 to 15.

SP is strictly a priority queue. When two PCs send 100% traffic to a same PC, all packets from the lower priority PC will be discarded. The default working mode is SP.

WRR is a weighted round robin queue, and under this mode, user can configure weight. If PCA and PCB sends 100% traffic to the same PCC, the PCC will receive packets from PCA and PCB according to the weight proportion in the corresponding queue.

WFQ is weighted fair queuing. Under this mode, user can configure guaranteed-rate and weight, and the guarantee is only available in WFQ mode. If PCA and PCB send 100% traffic to the same PCC, the PCC will receive packets from PCA and PCB according to the weight proportion and the guaranteed-rate in the corresponding queue.

```
admin@XorPlus# set class-of-service scheduler s1 mode WRR
admin@XorPlus# set class-of-service scheduler s1 weight 3
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

```
admin@XorPlus# set class-of-service scheduler s1 mode WFQ
admin@XorPlus# set class-of-service scheduler s1 weight 4
admin@XorPlus# set class-of-service scheduler s1 guaranteed-rate 8
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure a Forwarding Class

A forwarding class, which determines the queue number of the specified traffic type, should be configured after the scheduler when configuring QoS. The effective local-priority is 0 to 7.

```
admin@XorPlus# set class-of-service forwarding-class f1 local-priority 3
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure Scheduler Profile

Scheduler profile, which defines a queue used as a kind of scheduler on egress port, is the map of a forwarding class and a scheduler.

```
admin@XorPlus# set class-of-service scheduler-profile p1 forwarding-class f1 scheduler s1
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure Scheduler Profile to Specified Port

Scheduler profile should be configured to egress port. It only applies to egress packet. Ingress packet is invalid.

```
admin@XorPlus# set class-of-service interface ge-1/1/1 scheduler-profile p1
admin@XorPlus# commit
Commit OK.
Save done.
```

Configure a Classifier with IEEE 802.1/DSCP/ToS

A classifier should be configured first, which is used to specify the associated forwarding class. User can select a classifier trust mode, such as IEEE 802.1, DSCP, or ToS, according to need. It decides the priority trust model. Configure trust mode IEEE 802.1 as follows:

```
admin@XorPlus# set class-of-service classifier c1 trust-mode ieee-802.1
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure Classifier Relevant to Forwarding Class

After configuring a classifier trust mode, user can configure the classifier relevant to the specified forwarding class. Code point and scheduler should be configured at the same time. The code-point is matched with the forwarding class local-priority, meaning that when the flow matches the specified code point, the flow will enter the specified queue. When the classifier trust mode is IEEE 802.1 or ToS, the code point is 0 to 7. When the classifier trust mode is dscp, the code point is 0 to 63.

```
admin@XorPlus# set class-of-service classifier c1 forwarding-class f1 code-point 5
admin@XorPlus# set class-of-service classifier c1 forwarding-class f1 scheduler s1
admin@XorPlus# commit
Merging the configuration.
Commit OK.
Save done.
```

Configure Classifier to Specified Port

After configuring as above, the classifier should be applied to specified ports. It determines the port priority trust model, data stream and queue matching rules, scheduling model, weight, and guaranteed-rate. When the classifier configures the scheduler, the classifier should be used in the egress port. When the classifier configures code point, the classifier should be used in ingress port.

```
admin@XorPlus# set class-of-service interface ge-1/1/1 classifier c1
admin@XorPlus# commit
Waiting for merging configuration.
Commit OK.
Save done.
```

PFC Configuration Example

Priority Flow Control or PFC is a kind of flow control mechanism. The advantage of PFC over traditional Flow Control mechanisms is that PFC provides flow control based on per code point (priority). In other words, PFC provides a more granular form of flow control. This means that if traffic from one particular priority suffers from congestion, only that traffic is paused until congestion clears away, whereas traffic for other priorities continues unhindered. On each physical port, there are 8 (0 to 7) Class of Service (CoS) queues, if congestion is detected on the egress physical port, the ingress port will send a PAUSE frame to the transmitting node to pause transmission until the receiving node is ready to accept packets again. PFC applies only to packets entering a port.

PFC has a higher priority than flow control. So, for example, if both flow control and PFC are configured on the same port, PFC will have precedence over traditional flow control.

PFC uses the IEEE 802.1p CoS values in the IEEE 802.1Q VLAN tag to generate the flow control frame with corresponding priority on ingress physical port when egress physical port suffers congestion. It indicates the ingress port needs COS classifier configuration.

PFC configuration is applied using PFC profiles. The following is an example of PFC configuration.

Configure PFC Profile

PFC is disabled when drop value is set to true and enabled when drop value is set to false. The default value of drop is false.

For example as below, PFC is enabled on 013567 code-point by default, PFC is disabled on 24 code-point.

```
admin@XorPlus# set class-of-service pfc-profile pfc1 code-point 2 drop true
admin@XorPlus# set class-of-service pfc-profile pfc1 code-point 4 drop true
admin@XorPlus# commit
Commit OK.
Save done.
```

Apply PFC Profile to Port

```
admin@XorPlus# set class-of-service interface ge-1/1/1 pfc-profile pfc1
admin@XorPlus# commit
Commit OK.
Save done.
```

Show PFC Frame Statistics on Port

The class 0~7 in PFC frame corresponds to the following "802.1P" item. The value of "RxPFC" item will be incremented by 1 if ge-1/1/1 receives a PFC frame. The value of "TxPFC" item will be incremented by 1 if ge-1/1/1 sends out a PFC frame.

```
admin@XorPlus# run show class-of-service interface ge-1/1/1
Interface : ge-1/1/1
```

802.1P	Priority Flow Control	RxPFC	TxPFC
0	true	0	0
1	true	0	0
2	false	0	0
3	true	0	0
4	false	0	0
5	true	0	0
6	true	0	0
7	true	0	0

trust mode : no-trust
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0

Local-priority	Queue-Schedule	Code-points
0	SP,0kbps	
1	SP,0kbps	
2	SP,0kbps	
3	SP,0kbps	
4	SP,0kbps	
5	SP,0kbps	
6	SP,0kbps	
7	SP,0kbps	

Apply classifier based on IEEE 802.1P to ingress port

```

admin@XorPlus# set class-of-service classifier c1
admin@XorPlus# set class-of-service interface ge-1/1/1 classifier "c1"
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus# run show class-of-service interface ge-1/1/1
Interface : ge-1/1/1

```

802.1P	Priority Flow Control	RxPFC	TxPFC
0	true	0	0
1	true	0	0
2	false	0	0
3	true	0	0
4	false	0	0
5	true	0	0
6	true	0	0
7	true	0	0

trust mode : no-trust
Default ieee-802.1 : 0
Default dscp : 0
Default inet-precedence : 0

Local-priority	Queue-Schedule	Code-points
0	SP,0kbps	0
1	SP,0kbps	1
2	SP,0kbps	2
3	SP,0kbps	3
4	SP,0kbps	4
5	SP,0kbps	5
6	SP,0kbps	6
7	SP,0kbps	7



notice

Since version of 2.9.0, PFC has port limit. Different platform can support different port number with PFC .The configuration field as shown below :

Chip		Trident	Trident2	Trident2 plus	Helix4	Tomahawk(perxpe)
------	--	---------	----------	---------------	--------	------------------

1G	max PFC ports				18	
10G/25G	max PFC ports	32	42	56	11	12
40G	max PFC ports	19	25	33		7
100G	max PFC ports					3

In addition,if a 40G port configures PFC, the valid PFC port number of 10G/25G will reduce 2 .

OpenFlow in Crossflow Mode

OpenFlow is supported in PicOS, and user can configure PicOS switches with both legacy network protocol and Openflow, which will provide extreme flexibility in network deployment.

Configure ovssdb-server Locally

Check the ovssdb-server state on the switch. By default, the ovssdb-server listens to the local switch, and the ovs-vswitchd can only be configured by the local switch.

```
admin@XorPlus$ps aux|grep ovs
root      5174  0.5  0.7  46716  3900 ?        S<l  01:41   0:07
/xorplus/bin/system/tools/xorplus_monitor -c 80 -m 150 xorplus_lcmgr xorplus_sif
ovs-vswitchd
root      5205  0.0  0.1   6288   624 ?        Ss   01:42   0:00 ovssdb-server: monitoring
pid 5206 (healthy)

root      5206  0.0  0.4   6432  2124 ?        S    01:42   0:00 /ovs/sbin/ovssdb-server
--pidfile --log-file --detach --monitor --remote=punix:/ovs/var/run/openvswitch/db.sock
--remote=db:Open_vSwitch,Manager,target --private-key=db:hardware_vtep,SSL,private_key
--certificate=db:hardware_vtep,SSL,certificate
--bootstrap-ca-cert=db:hardware_vtep,SSL,ca_cert /ovs/ovs-vswitchd.conf.db
/xorplus/config/vtep.db
root      5208  0.0  1.0  49580  5388 ?        Sl   01:42   0:00 /ovs/sbin/ovs-vswitchd
--enable-shared-lcmgr
admin    5219  0.0  0.1   2128   684 ttyS0    S+   02:01   0:00 grep --color=auto ovs
```

Configure ovs-vswitchd Locally

```
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
admin@XorPlus$ovs-vsctl add-port br0 ge-1/1/21 -- set interface ge-1/1/21 type=pica8
admin@XorPlus$ovs-vsctl set-controller br0 tcp:10.10.51.59:6633
```

Configure ovssdb-server Remotely

Modify the ovssdb-server state on the switch. Configure the ovssdb-server listening switch management-ethernet interface IP 10.10.51.138 and PORT 6640.

```
admin@XorPlus$ps aux|grep ovs
root      5174  0.6  0.7  46716  3900 ?        S<l  01:41   0:10
/xorplus/bin/system/tools/xorplus_monitor -c 80 -m 150 xorplus_lcmgr xorplus_sif
ovs-vswitchd
root      5205  0.0  0.1   6288   624 ?        Ss   01:42   0:00 ovssdb-server: monitoring
pid 5206 (healthy)

root      5206  0.0  0.4   6432  2136 ?        S    01:42   0:00 /ovs/sbin/ovssdb-server
--pidfile --log-file --detach --monitor --remote=punix:/ovs/var/run/openvswitch/db.sock
--remote=db:Open_vSwitch,Manager,target --private-key=db:hardware_vtep,SSL,private_key
--certificate=db:hardware_vtep,SSL,certificate
--bootstrap-ca-cert=db:hardware_vtep,SSL,ca_cert /ovs/ovs-vswitchd.conf.db
/xorplus/config/vtep.db
root      5208  0.3  1.2  75456  6520 ?        Sl   01:42   0:05 /ovs/sbin/ovs-vswitchd
--enable-shared-lcmgr
admin    5262  0.0  0.1   2128   684 ttyS0    S+   02:09   0:00 grep --color=auto ovs
```

```

admin@XorPlus$
admin@XorPlus$ sudo kill -9 5206
admin@XorPlus$
admin@XorPlus$ sudo /ovs/sbin/ovsdb-server --pidfile --log-file --detach --monitor
or --remote=ptcp:6640:10.10.51.144 --remote=punix:/ovs/var/run/openvswitch/db.sock
--remote=db:Open_vSwitch,Manager,target --private-key=db:hardware_vtep,SSL,
private_key --certificate=db:hardware_vtep,SSL,certificate --bootstrap-ca-cert=
db:hardware_vtep,SSL,ca_cert /ovs/ovs-vswitchd.conf.db /xorplus/config/vtep.db
2002-06-25T02:15:31Z|00001|vlog|INFO|opened log file
/ovs/var/log/openvswitch/ovsdb-server.log
admin@XorPlus$
admin@XorPlus$ ps aux|grep ovs
root      5174  0.5  0.7 46716 3900 ?        S<l  01:41   0:12
/xorplus/bin/system/tools/xorplus_monitor -c 80 -m 150 xorplus_lcmgr xorplus_sif
ovs-vswitchd
root      5208  0.2  1.2 75456 6520 ?        Sl   01:42   0:05 /ovs/sbin/ovs-vswitchd
--enable-shared-lcmgr
root      5267  0.0  0.1  6288  628 ?        Ss   02:15   0:00 ovsdb-server: monitoring
pid 5268 (healthy)

root      5268  0.5  0.4  6432 2124 ?        S    02:15   0:00 /ovs/sbin/ovsdb-server
--pidfile --log-file --detach --monitor --remote=ptcp:6640:10.10.51.144
--remote=punix:/ovs/var/run/openvswitch/db.sock --remote=db:Open_vSwitch,Manager,target
--private-key=db:hardware_vtep,SSL,private_key
--certificate=db:hardware_vtep,SSL,certificate
--bootstrap-ca-cert=db:hardware_vtep,SSL,ca_cert /ovs/ovs-vswitchd.conf.db
/xorplus/config/vtep.db
admin     5270  0.0  0.1  2128  684 ttyS0    S+   02:15   0:00 grep --color=auto ovs
admin@XorPlus$

```

Configure ovs-vswitchd by Remote Server 10.10.50.42

```

root@dev-42:~# ovs-vsctl --db=tcpx:10.10.51.144:6640 add-br br0 -- set bridge br0
datapath_type=pica8
root@dev-42:~# ovs-vsctl --db=tcpx:10.10.51.144:6640 add-port br0 ge-1/1/21 -- set interface
ge-1/1/21 type=pica8
root@dev-42:~# ovs-vsctl --db=tcpx:10.10.51.144:6640 set-controller br0 tcpx:10.10.51.59:6633

```

- Crossflow Mode Introduction
- CrossFlow Mode Known Limitations
- Configurations Example in Crossflow Mode

Crossflow Mode Introduction

This chapter describes the details of OpenFlow via CrossFlow mode. CrossFlow mode allows traditional L2/L3 and OpenFlow protocols to run simultaneously on the same physical switch.



CrossFlow mode has been greatly improved starting in PicOS 2.6.2. Now, most of the OVS mode commands are available in CrossFlow mode.

OVS features that are available in the CrossFlow mode:

- controller (test ok)
- egress-mode (test ok)
- combined-mode (test ok)
- multi-table (test ok)
- max-ecmp-ports (test ok)
- loopback-enable (test ok)
- flow-counter-mode (test ok)
- GRE (test ok)
- meter (test ok)
- group (test ok)

- Q-in-Q (test ok)
- udf-mode(test ok)
- manager(test ok)
- mpls flow entries(test ok)

OVS features that need to be configured in Xorplus mode:

- qos (test ok)
- qe-port-mode (test ok)
- lag-advance-hash-mapping-fields (test ok)
- LACP (test ok)

The OVS features which are unavailable in CrossFlow mode:

- proxy-arp
- proxy-icmpv6
- pbb flow entries
- VxLAN
- L2MPLS
- L2GRE
- match-mode



Configuration Notes:

- Only Openflow ports can be added to an Openflow lag as members, and only Crossflow ports can be added to a Crossflow lag as members. All lag's members are configured in Switching CLI.
- All LACP packets are handled by Switching.
- Besides LACP, all other protocol packets received from Openflow ports will be sent to OVS. For the packets received from Crossflow ports, user can configure either Switching or OVS to handle the packets. By default, these packets are handled by Switching.
- All port speed, mtu, VLAN members, and other information is configured in Switching CLI as usual, and these configurations in OVS CLI will not take effect.
- When add flow with matching mpls, mpls_label cannot be configured as 0-15.

The Switch Hardware Architecture page describes the meaning of some terms used here, such as TCAM or FIB.

CrossFlow Mode Introduction

In CrossFlow mode, switches can achieve most of the functions that exist in OVS mode, including: basic flow function, meter, group, multi-table, Q-in Q and so on.

Like in PicOS OVS mode, in CrossFlow mode, OpenFlow 1.0, OpenFlow 1.1, OpenFlow 1.2, OpenFlow 1.3, and OpenFlow 1.4 are supported. User can configure any supported version in the CLI.

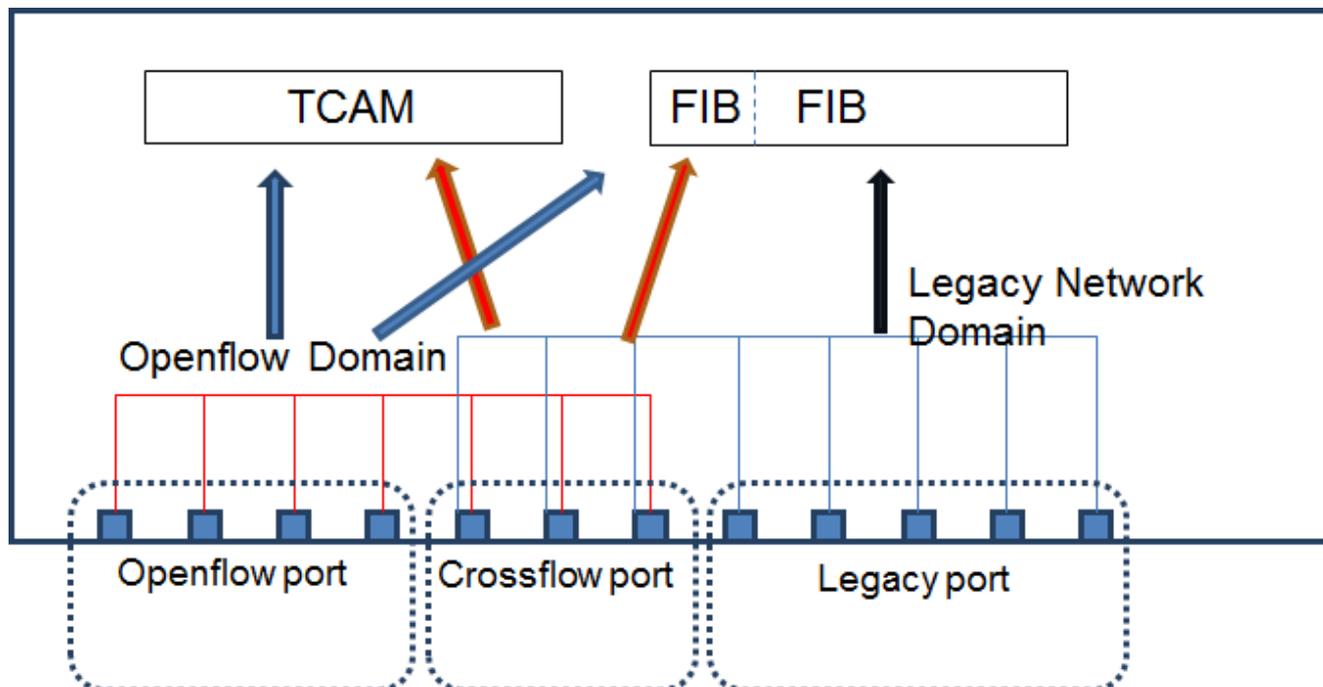
After PicOS 2.6.2, ports in the switch can be configured in any one of these three modes: Legacy, Crossflow or Openflow. The difference in configuration is:

Legacy port: disable crossflow

Crossflow port: enable crossflow and enable local-control

Openflow port: enable crossflow and disable local-control

After PicOS 2.4, Openflow can also support the multiple-table control, which means that Openflow can configure some flows in FIB (mac address table and routing table). The flow should match some conditions to install them in FIB. Please refer to the manual in the PicOS Openflow Configuration documentation.



As the figure above shows, the switch ports working status in Crossflow mode is shown. The Openflow flow can look up two tables including TCAM and FIB (multiple tables). The ports in Crossflow mode can be summarized as follows:

Enable Crossflow mode and local-control disable (Openflow port):

1. The port is completely controlled by Openflow
2. All broadcast turned off and auto learning turned off
3. Packet forwarded by looking up the TCAM table default
4. Packet forwarded by looking up the FIB once enable multi-table and configure L2/L3 flow table

Enable Crossflow mode and local-control enable (Crossflow port):

1. The port is controlled by both local legacy stack and Openflow.
2. All broadcast is turned on and auto learning is turned on
3. Packet is forwarded by looking up the FIB (FDB/routing table) and TCAM table
4. No matter if multi-table is enabled or not, packets all can forward by looking up the FIB

Disable Crossflow mode (Legacy port):

1. The port is controlled by local legacy stacks.
2. All broadcast is turned on and auto learning is turned on.
3. Packet is forwarded by looking up the FIB (FDB/routing table).

In Crossflow mode, user can enable L2/L3 mode as a PicOS OVS multi-table function. If the L2/L3 mode is enabled, the FIB table resource will be shared by legacy ports and Crossflow ports, and data traffic cannot mix between the different type ports. User can allocate the resource for special ports as the following command shows.

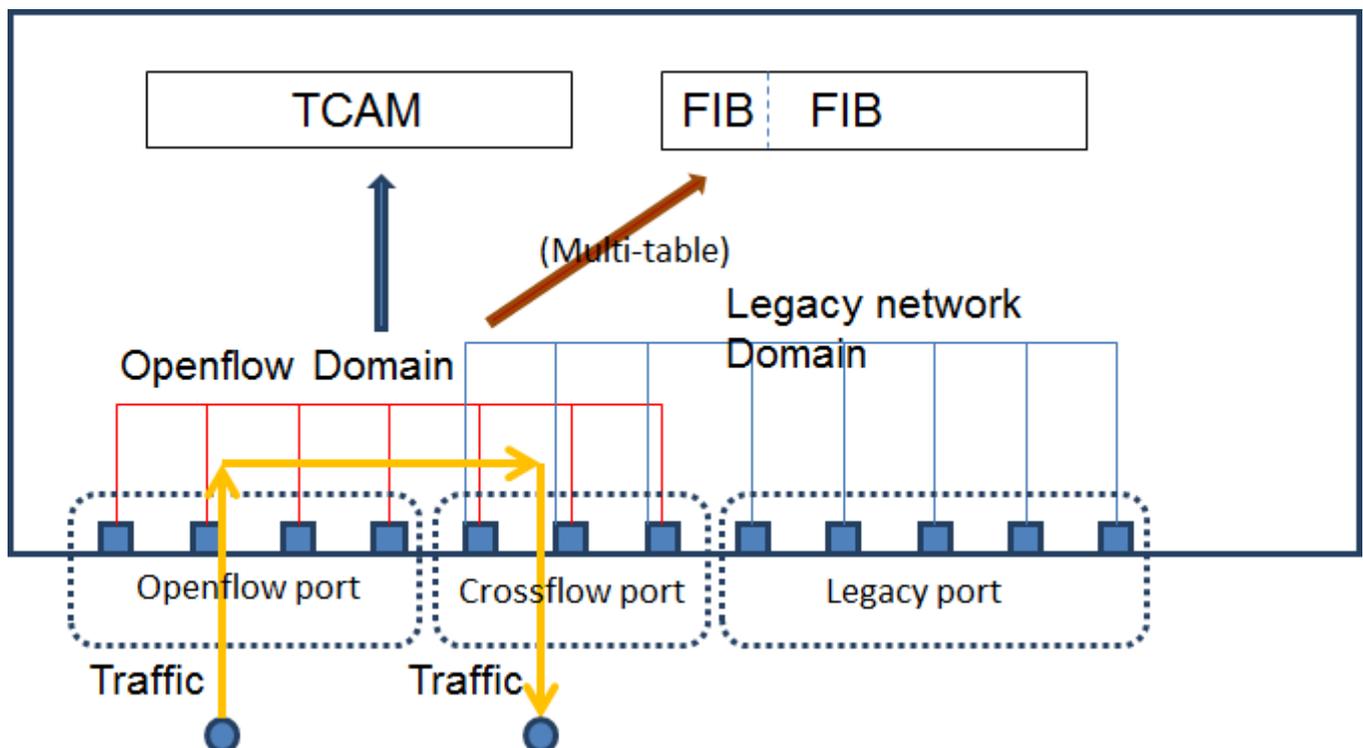
```
set interface stm firewall-table ingress 100
set interface stm mac-table 20000
set interface stm ipv4-route 6000
set interface stm ipv6-route 1000
```

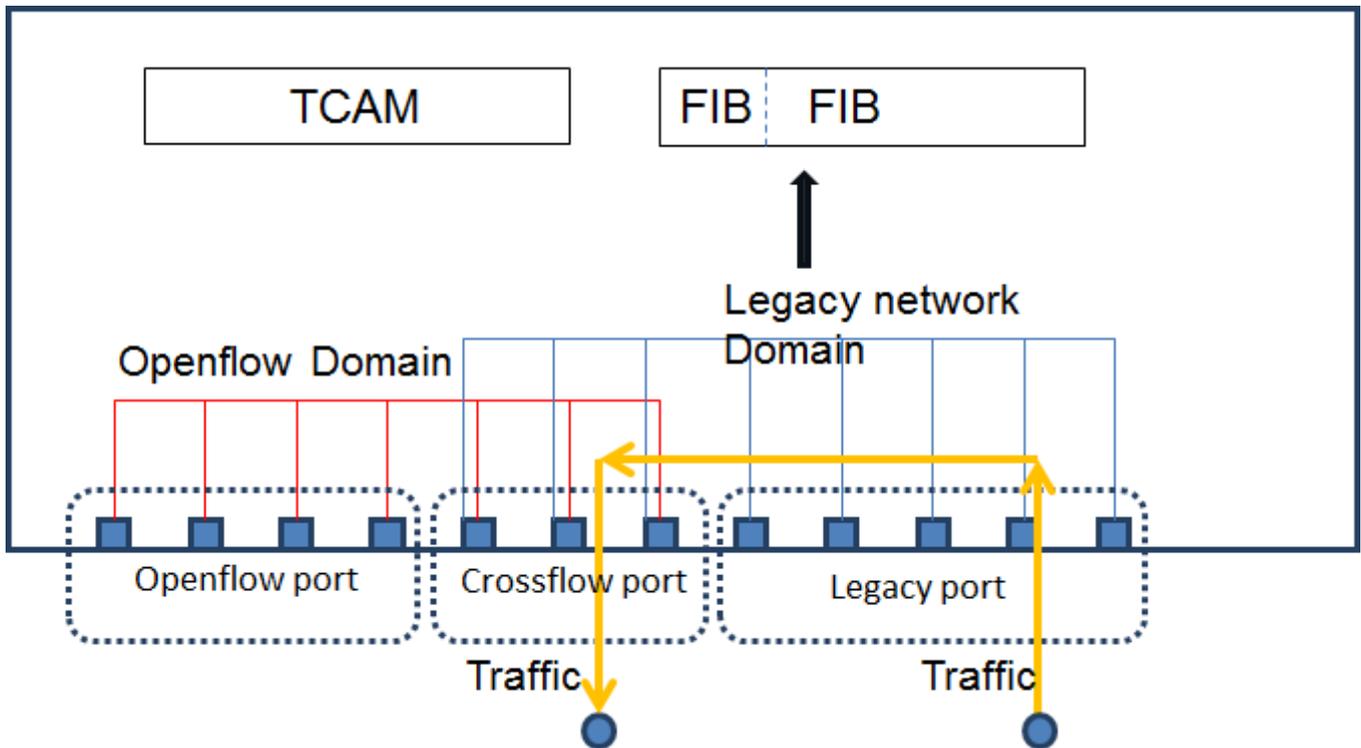
The first command, **set interface stm firewall-table ingress 100**, means allocate the TCAM resource for ingress ACL rule, and part of the rest of resource for Crossflow port and Openflow port;

The second command, **set interface stm mac-table 20000**, means allocate the FDB table resource for legacy port, and the rest of resource for Crossflow port and Openflow port.

The last two commands, **set interface stm ipv4-route 6000**, and **set interface stm ipv6-route 1000**, means allocate the route table resource for legacy port, and the rest of resources for Crossflow port and Openflow port.

In Crossflow mode, traffic can only be forwarded in the OpenFlow or Legacy network domain (as shown in the following figures). The traffic between Openflow port and Crossflow port forwarded by Tcam flow table default, if user wants to forward by FIB, enable multi-table and configure L2/L3 flow entries. The traffic between Crossflow port and Legacy port is forwarded by FIB.





Install a flow with action "NORMAL":

1. By default, the traffic output port will be replaced by the TCAM flow's action after packets are switched and routed. e.g. A packet matches a route and decides the output port (ge-1/1/10) by a route, then it also matches a TCAM flow which has the output port ge-1/1/12, Then, the packet will be forwarded to ge-1/1/12 because the TCAM flow action will replace the original route output port.
2. When a flow in TCAM has an action "NORMAL", the packet which has decided the output port after route table, will still be forwarded to route output port. "NORMAL" means there is no change to the packet output port.
3. When a route packet will not match any flows in TCAM, the output port also will not be changed.

CrossFlow Mode Known Limitations

Conflict Management between OpenFlow and CLI Configurations

The physical configurations of the FDB table and routing table of Openflow should not conflict with the CLI configurations. This means that Openflow can install a flow in the routing table at the same time that legacy network installs a flow in routing table without conflicts. If these configurations conflict, the later configuration will fail.

Firewall filters and OpenFlow TCAM Flows may conflict. For instance, the OpenFlow rule could be dropping a specific type of packet when the Firewall filters are forwarding them. In this case, both rules are performed concurrently and independently, and the results of the matches are merged. When there are no conflicting results, all results are applied. When there are overlapping and conflicting results, the conflicted parts of the result are selected based on priorities. DROP action has the highest priority, followed by REDIRECT, REPLACE, and TCAM slice number. (OpenFlow rules and Firewall rules are placed in different TCAM Slices.)

In CrossFlow mode, the ACL filter has a higher priority than that of Openflow flow entry.

Because arp works in different group with ip, so if there is one or more openflow ports, the arp packets will be dropped because drop actions has the highest priority. And you can add a flow like this: `ovs-ofctl add-flow br0 priority=1,actions=normal`

Default Drop In OpenFlow

There is a default drop flow in the system when user enables Crossflow mode, and this drop flow only applies to the Openflow ports.

Other Limitations

From PicOS-2.6, some ports can work within both the legacy network domain and the Openflow domain. We call these ports Crossflow ports.

If user enables multi-table and configures L2/L3 flow entry on the Openflow port, traffic from the legacy port can also match the L2/L3 flow entry and be forward on the Openflow port.

If packets can match TCAM flow entry and route table at the same time, the TCAM flow entry has the higher priority. Because the packets must go through FIB table, if there is no mod-src-mac in the TCAM flow entry, the packets will be modified by the src-mac in FIB table then go out as a TCAM flow entry.

If a packet needs match to a Crossflow port, it must have `in_port` in match field when adding flow.

Configurations Example in Crossflow Mode

Configure the OpenFlow Port in CrossFlow Mode

Commands:

```
admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 crossflow local-control false
admin@XorPlus# set vlans vlan-id 2,2000,4094
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan members 2000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching vlan
```

```
members 4094
admin@XorPlus# commit
```

Commands in Linux:

```
admin@XorPlus$ovs-vsctl list pica8
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
admin@XorPlus$ovs-vsctl add-port br0 ge-1/1/1 -- set Interface ge-1/1/1 type=pica8
```

Configure the Hybrid Port in CrossFlow Mode

Commands:

```
admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 crossflow enable true
admin@XorPlus# set vlans vlan-id 2,2000,4094
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 2
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 2000
admin@XorPlus# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching vlan
members 4094
admin@XorPlus# commit
```

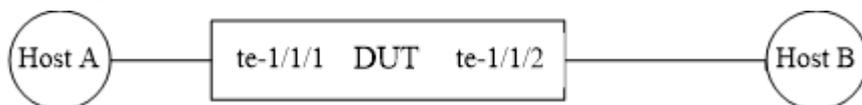
Commands in Linux:

```
admin@XorPlus$ovs-vsctl list pica8
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
admin@XorPlus$ovs-vsctl add-port br0 ge-1/1/2 -- set Interface ge-1/1/2 type=crossflow
```

Examples

Basic Configurations

topology



Step 1: Configure port te-1/1/1 as OpenFlow port and te-1/1/2 as CrossFlow port

```
admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow local-control false
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow enable true
admin@XorPlus# set vlans vlan-id 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching port-mode
```

```
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 speed 1000
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 speed 1000
admin@XorPlus# commit
```

Step 2: Exit the Xorplus system, then enter Linux system

```
admin@XorPlus#exit
admin@XorPlus>exit
admin@XorPlus$
```

Step 3: Create a new bridge named br0

```
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
```

Step 4: Add ports to br0

```
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/1 -- set Interface te-1/1/1 type=pica8
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/2 -- set Interface te-1/1/2 type=crossflow
```

Step 5: Add a flow

```
admin@XorPlus$ovs-ofctl add-flow br0 in_port=1,actions=output:2
```

Step 6: Send packets to te-1/1/1

Send untagged packets to te-1/1/1 that match this flow. Then, te-1/1/2 will forward the packets (with no vlan). Send packets with vlan 2 to te-1/1/1. Then, te-1/1/2 will forward the packets (with vlan 2).

Lag Configurations

Step 1: Set lag interface (Only Openflow ports can be added to an openflow lag as members, and only Crossflow ports can be added to a Crossflow lag as members.)

```
admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow local-control false
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow enable true
admin@XorPlus# commit
admin@XorPlus#set interface aggregate-ethernet ae1 crossflow enable true
admin@XorPlus#set interface aggregate-ethernet ae1 crossflow local-control false
admin@XorPlus#set interface aggregate-ethernet ae2 crossflow enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 speed 1000
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 speed 1000
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 ether-options 802.3ad ae1
```

```

admin@XorPlus# set interface gigabit-ethernet te-1/1/2 ether-options 802.3ad ae2
admin@XorPlus# commit
admin@XorPlus#set vlans vlan-id 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae1 family ethernet-switching vlan members
2,10,20
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface aggregate-ethernet ae2 family ethernet-switching vlan members
2,10,20
admin@XorPlus# commit

```

Step 2: Exit the Xorplus system, then enter Linux system

```

admin@XorPlus#exit
admin@XorPlus>exit
admin@XorPlus$

```

Step 3: Create a new bridge named br0.

```

admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8

```

Step 4: Add ports to br0.

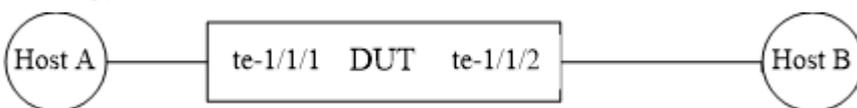
```

admin@XorPlus$ovs-vsctl add-port br0 ae1 -- set Interface ae1 type=pica8_lag
admin@XorPlus$ovs-vsctl add-port br0 ae2 -- set Interface ae2 type=pica8_lag

```

Flow Priority Configurations

topology



Step 1: Configure two ports as Openflow ports

```

admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow local-control false
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow local-control false
admin@XorPlus# set vlans vlan-id 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching port-mode
trunk

```

```
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# commit
```

Step 2: Exit the Xorplus system, then enter Linux system

```
admin@XorPlus#exit
admin@XorPlus>exit
admin@XorPlus$
```

Step 3: Create a new bridge named br0.

```
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
```

Step 4: Add ports to br0.

```
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/1 -- set Interface te-1/1/1 type=pica8
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/2 -- set Interface te-1/1/2 type=pica8
```

Step 5: Add two flows

```
admin@XorPlus$ovs-ofctl add-flow br0 in_port=1,dl_src=22:11:11:11:11:11,actions=output:2
admin@XorPlus$ovs-ofctl add-flow br0
in_port=1,priority=50000,dl_src=22:11:11:11:11:11,actions=mod_dl_src:22:22:22:22:22:22,outp
ut:2
```

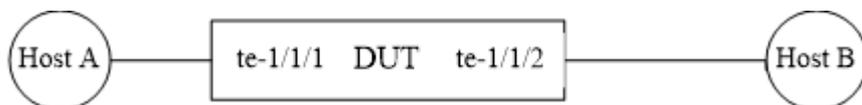
Step 6: Send packets to te-1/1/1

Send untagged packets to te-1/1/1 that match this flow, then te-1/1/2 will forward the packets (with no vlan), and the packets' source mac address is modified to 22:22:22:22:22:22, because the priority of the second flow is higher than that of the first flow.

Send packets with vlan 2 to te-1/1/1, then te-1/1/2 will forward the packets with vlan 2 and the packets' source mac address is modified to 22:22:22:22:22:22.

FDB Configurations

topology



Step 1: Configure two ports as Openflow ports

Step 2: Exit the Xorplus system, then enter Linux system

```
admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow local-control false
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow local-control false
admin@XorPlus#set vlans vlan-id 2,10,20
```

```

admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching port-mode
trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching vlan
members 2,10,20
admin@XorPlus# commit
admin@XorPlus#exit
admin@XorPlus>exit
admin@XorPlus$

```

Step 3: Create a new bridge named br0.

```

admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8

```

Step 4: Add ports to br0.

```

admin@XorPlus$ovs-vsctl add-port br0 te-1/1/1 -- set Interface te-1/1/1 type=pica8
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/2 -- set Interface te-1/1/2 type=pica8

```

Step 5: Set table 1 to FDB table

```

admin@XorPlus$ovs-vsctl set-l2-mode true 1

```

Step 6: Add a flow

```

admin@XorPlus$ovs-ofctl add-flow br0
table=1,dl_dst=22:22:22:22:22:22,dl_vlan=10,actions=output:2

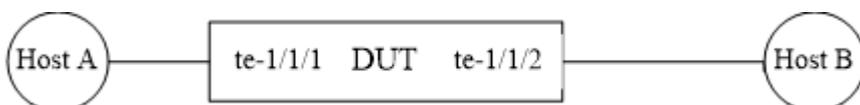
```

Flows must match `dl_dst`, `dl_vlan` and output port if they want to be stored in FDB table. Table number of FDB table is 251 by default. User can specify another table as the FDB table instead of the 251 by using this command: ***ovs-vsctl set-l2-mode true [table number]***.

If you want flows to be stored in ROUTE table, flows must match `dl_dst`, `dl_vlan`, `dl_type`, `nw_dst`, and `mod_dl_dst` in action, and the default table number of ROUTE is 252. Use command ***ovs-vsctl set-l3-mode true [table number]*** to set route table.

Route Configurations

topology



Step 1: Configure two ports as Openflow ports

Step 2: Exit the Xorplus system, then enter Linux system

```

admin@XorPlus# set xovs enable true
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 crossflow local-control false
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow enable true
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 crossflow local-control false
admin@XorPlus# set vlans vlan-id 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/1 family ethernet-switching vlan members 2,10,20
admin@XorPlus# commit
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching port-mode trunk
admin@XorPlus# set interface gigabit-ethernet te-1/1/2 family ethernet-switching vlan members 2,10,20
admin@XorPlus# commit
admin@XorPlus# exit
admin@XorPlus>exit
admin@XorPlus$

```

Step 3: Create a new bridge named br0.

```
admin@XorPlus$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
```

Step 4: Add ports to br0.

```
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/1 -- set Interface te-1/1/1 type=pica8
admin@XorPlus$ovs-vsctl add-port br0 te-1/1/2 -- set Interface te-1/1/2 type=pica8
```

Step 5: enable L3-mode and the default table is 252

```
admin@XorPlus$ovs-vsctl set-l3-mode true
```

Step 6: Add a route flow

```
admin@XorPlus$ovs-ofctl add-flow br0
table=252,dl_vlan=1,dl_dst=22:00:00:00:00:00,ip,nw_dst=1.1.1.100/24,actions=set_field:2-\>vlan_vid,set_field:22:22:22:22:22:22-\>eth_dst,output:2
```

VXLAN Configuration

- VXLAN Configuration Guide
- VXLAN Base Configuration Example
- VXLAN ECMP Configuration
- OVSDB VTEP Configuration
 - Configuring an OVSDB VTEP
 - OVSDB VTEP with Midonet Configuration
 - OVSDB VTEP with NSX Configuration
 - OVSDB VTEP with vtep-ctl Configuration Examples

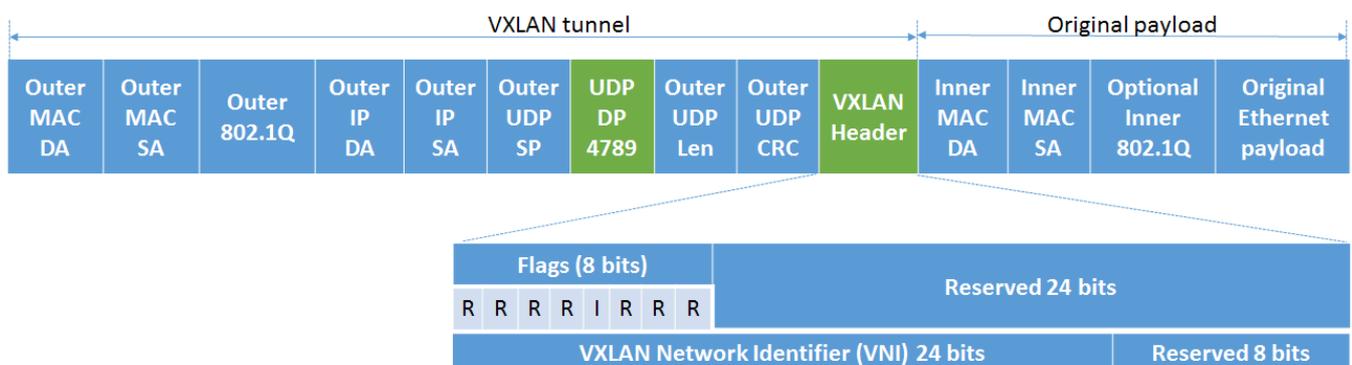
About VXLAN

Virtual Extensible LAN (VXLAN) is an overlay network virtualization technology. An overlay network is a virtual network that is built on top of existing network Layer 2 and Layer 3 technologies to support elastic compute architectures. VXLAN makes it easier for network engineers to scale out a cloud computing environment while logically isolating cloud apps and tenants.

VXLAN Technology

VXLAN uses UDP-based encapsulation to tunnel Ethernet frames and transfers original data packets as tunnel payloads. With the outer UDP tunnel, inner payload data can be quickly transferred on the layer 2 and layer 3 networks. To provide the capability of broadcast domain addressing, the VXLAN technology uses layer 3 IP multicast to replace the Ethernet broadcast. Therefore, the broadcast, unknown unicast, and multicast (BUM) packets can be transferred on virtual networks through broadcasting. For more VXLAN details, please read RFC7348.

VXLAN Standards



As shown in Figure 1-1, a VXLAN packet consists of the outer encapsulation and the inner payloads.

- Flags (8 bits): The flag I must be set to 1 for a valid VXLAN Network Identifier (VNI). The other 7 bits (labeled as R) are reserved fields and must be set to 0 on transmit and ignored on receive.

- VXLAN segment ID or VXLAN VNI: This parameter contains 24 bits and is used to designate the individual VXLAN overlay network on which the VMs are located.
- Reserved fields (24 bits and 8 bits): This parameter must be set to 0 on transmit and ignored on receive.
- The destination port number assigned to the outer tunnel is 4789, which is dedicated.

 **The switch platforms which use this feature are:**

- Trident2
- Trident2+
- Tomahawk

VXLAN Inner 802.1Q

Encapsulation mode

Encapsulation means the flow from access ports to network ports. Use one of the following options to specify actions about 802.1Q tag while encapsulation.

- **none**: Nothing will change, untagged packets will stay untagged, tagged packets will stay tagged.
- **service-vlan-add**: Add 802.1Q tag for untagged packets, and nothing changed with tagged packets. Encapsulation vlan is required.
- **service-vlan-add-delete**: Add 802.1Q tag for untagged packets, and delete tag for tagged packets. Encapsulation vlan is required.
- **service-vlan-add-replace**: Add 802.1Q tag for untagged packets, and replace tag for tagged packets. Encapsulation vlan is required.
- **service-vlan-delete**: Delete 802.1Q tag for tagged packets, and nothing changed with untagged packets. This is default value according to RFC 7348.
- **service-vlan-replace**: Replace vlan id of 802.1Q tag for tagged packets, and nothing changed with untagged packets. Encapsulation vlan is required.

Decapsulation-mode

Decapsulation means the flow from network ports to access ports.

- **none**: Nothing will change, untagged packets will stay untagged, tagged packets will stay tagged.
- **service-vlan-add**: From network ports to access ports, add 802.1Q tag for both untagged/tagged packets. If the access port is matched by port and vlan, the vlan id of the tag being added will be that vlan, otherwise will be PVID of that port.
- **service-vlan-add-delete**: From network ports to access ports, add 802.1Q tag for both untagged/tagged packets. If the access port is matched by port and vlan, the vlan id of the tag being added will be that vlan, otherwise will be PVID of that port. From access to access, delete tag for tagged packets.

- **service-vlan-add-replace:** From network ports to access ports, add 802.1Q tag for both untagged/tagged packets. If the access port is matched by port and vlan, the vlan id of the tag being added will be that vlan, otherwise will be PVID of that port. From access to access, replace tag for tagged packets. This is the default value.
- **service-vlan-delete:** From access to access, delete tag for tagged packets.
- **service-vlan-replace:** From access to access, replace tag for tagged packets.

Based on the above description, please see the following three tables for the detailed traffic changes.

The below table shows the traffic changes in the case that interfaces in the access side are binded with a vxlan in the network side.

	Access->Access	Access->Network	Network->Access
none	untag-->untag tag-->remain tag	untag->untag tag->remain tag	untag-->untag tag-->remain tag
service-vlan-add	untag->tag(add pvid) tag->remain tag	untag->add encapsulation vlan tag tag->remain tag	untag-->tag(add pvid) tag-->double tag(outer layer add pvid)
service-vlan-add-delete	untag->tag(add pvid) tag-->untag	untag-->add encapsulation vlan tag->untag(being deleted)	untag-->tag(add pvid) tag-->double tag(outer layer add pvid)
service-vlan-add-replace	untag->tag(add pvid) tag-->new tag(replaced with pvid)	untag-->add encapsulation vlan tag->tag(changed to encapsulation vlan)	untag-->tag(add pvid) tag-->double tag(outer layer add pvid)
service-vlan-delete	untag-->untag tag-->untag	untag->untag tag->untag	untag-->untag tag-->remain tag
service-vlan-replace	untag-->untag tag-->new tag(replaced with pvid)	untag->untag tag->tag(changed to encapsulation vlan)	untag-->untag tag-->remain tag

The below table shows the traffic changes in the case that the interfaces and vlans in the access side are binded with a vxlan in the network side.

	Access->Access	Access->Network	Network->Access
none	tag-->remain tag	tag->remain tag	untag-->untag tag-->remain tag
service-vlan-add	tag->remain tag	tag->remain tag	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)
service-vlan-add-delete	tag-->untag	tag->untag(being deleted)	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)
service-vlan-add-replace	tag-->new tag(replaced with vxlan-vlan)	tag->tag(changed to encapsulation vlan)	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)

	Access->Access	Access->Network	Network->Access
service-vlan-delete	tag-->untag	tag->untag	untag-->untag tag-->remain tag
service-vlan-replace	tag-->new tag(replaced with vxlan-vlan)	tag->tag(changed to encapsulation vlan)	untag-->untag tag-->remain tag

The below table shows the traffic changes in the case that only vlans in the access side are binded with a vxlan in the network side.

	Access->Access	Access->Network	Network->Access
none	tag-->remain tag	tag->remain tag	untag-->untag tag-->remain tag
service-vlan-add	tag->remain tag	tag->remain tag	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)
service-vlan-add-delete	tag-->untag	tag->untag(being deleted)	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)
service-vlan-add-replace	tag-->remain tag	tag->tag(changed to encapsulation vlan)	untag-->tag(add vxlan-vlan) tag-->double tag(outer layer add vxlan-vlan)
service-vlan-delete	tag-->untag	tag->untag	untag-->untag tag-->remain tag
service-vlan-replace	tag-->remain tag	tag->tag(changed to encapsulation vlan)	untag-->untag tag-->remain tag

VXLAN ECMP

In L2/L3, VXLAN ECMP is supported. Picos supports up to 32-way ECMP.

- The VXLAN ECMP does not need special configuration. It entirely depends on the routing ECMP. The route ECMP configure link: [Configuring ECMP \(Equal-Cost Multipath Routing\)](#)
- PicOS uses info from VXLAN header for hash calculation to ensure better performance.

VXLAN Mac Learning

The VTEP performs source MAC learning on the VNI as a Layer 2 switch.

- The switch receives traffic from the local VTEP to the remote VTEP, the VTEP learns the source MAC address in the access port.
- The switch receives traffic from the remote VTEP to the local VTEP, the VTEP learns the source MAC address in the network port.

A VNI MAC address table includes the following types of MAC address entries:

- Access port--Dynamic MAC address entries learned from the local VTEP. VXLAN does not support local configure static MAC address.
- Network port--Include static and dynamic MAC entries.

Static mac--Configure static mac address entries on VXLAN tunnel interfaces.

Dynamic mac--The MAC address entries learned from incoming traffic on VXLAN tunnels. The learned MAC addresses are contained in the inner Ethernet header source MAC.

On network port, the configure static mac entry has higher priority than dynamic mac entries.

VXLAN Traffic Forwarding

Unicast Traffic

- The switch receives traffic from the access port. The VTEP encapsulates the original Ethernet frame with an outer MAC header, outer IP header, and a VXLAN header. The source IP address is the source VTEP's VXLAN tunnel source IP address.
- The local VTEP forwards the encapsulates packets to the VXLAN tunnel a destination IP address.
- The remote VTEP decapsulates the packet and forwards the frame to access port.

Broadcast and Unknown Traffic

- The switch receives traffic from the access port. The VTEP encapsulates the original Ethernet frame with an outer MAC header, outer IP header and a VXLAN header. The source IP address is the source VTEP's VXLAN tunnel source IP address.
- The local VTEP flood encapsulates packets to the VXLAN tunnel all destination IP address.
- The all remote VTEP decapsulates the packet and forwards the frame to access port.

Configure map port to VXLAN VNI or map port&vlan to VXLAN VNI

Step

VxLAN supported on PicOS L2/L3 switch, the commands configure step as below.

- Configure VXLAN source interface

```
set vxlans source-interface loopback address 10.10.10.25
commit
```

- Create VXLAN VNI

```
set vxlans vni 10010
commit
```

- Configure vtep address for VXLAN VNI

```
set vxlans vni 10010 flood vtep 10.10.10.12
commit
```

- Add VXLAN port into VXLAN VNI

```
set vxlans vni 10010 interface te-1/1/40 vlan 100
commit
```

Configure map vlan to VXLAN VNI Step

VxLAN supported on PicOS L2/L3 switch. To configure Step, please see below.

- Configure VXLAN source interface

```
set vxlans source-interface loopback address 10.10.10.25
commit
```

- Create VXLAN VNI

```
set vxlans vni 10010
commit
```

- Configure vtep address for VXLAN VNI

```
set vxlans vni 10010 flood vtep 10.10.10.12
commit
```

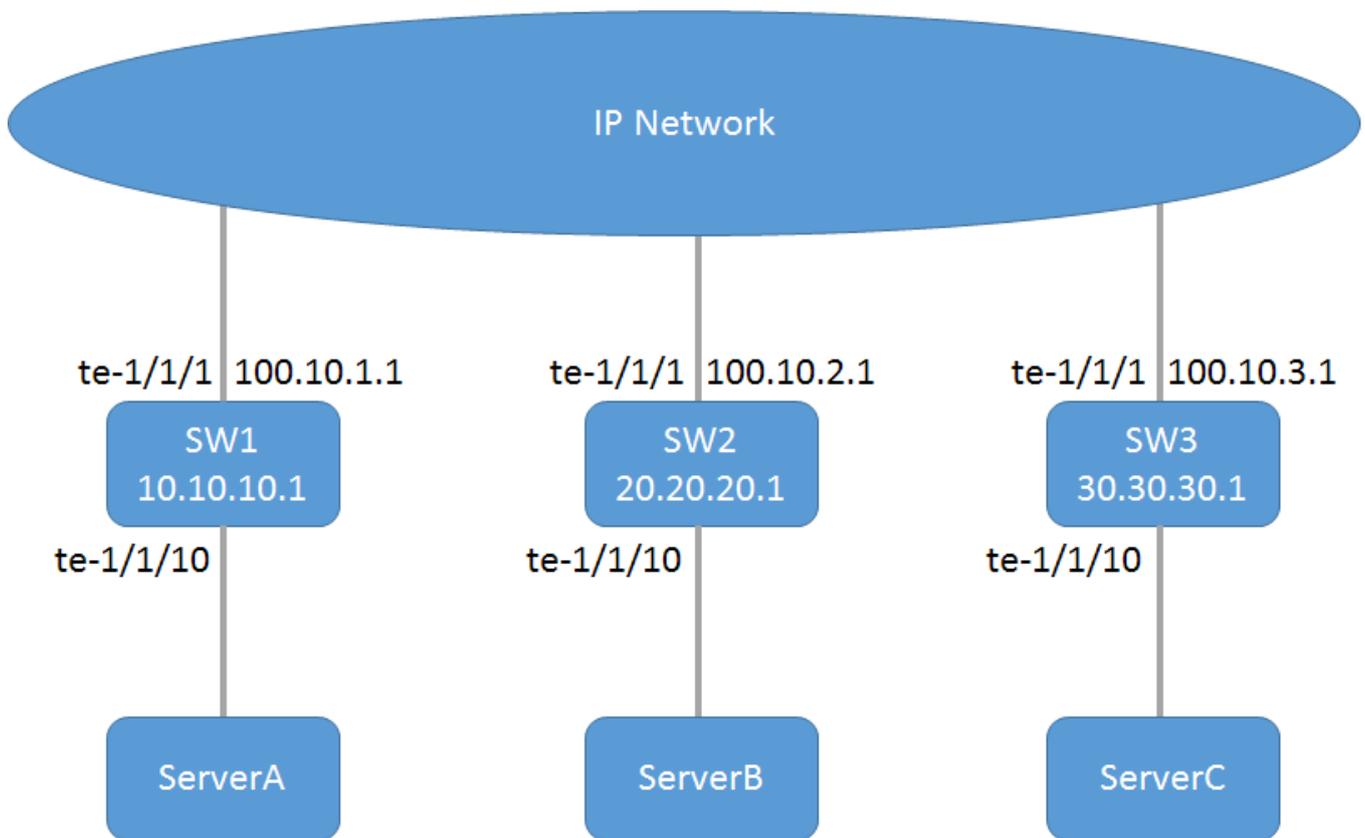
- Enable VXLAN VNI map with vlan

```
set vxlans vni-map-vlan true
commit
```

- Add vlan into VXLAN VNI

```
set vxlans vni 10010 vlan 100
commit
```

Topology



SW1 Configure

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
set vlan-interface loopback address 10.10.10.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.1.1 prefix-length 24
set protocols static route 20.20.20.1/32 next-hop 100.10.1.2
set protocols static route 30.30.30.1/32 next-hop 100.10.1.3
set vxlans source-interface loopback address 10.10.10.1
set vxlans vni 1000 interface te-1/1/10 vlan 10
set vxlans vni 1000 flood vtep 20.20.20.1
set vxlans vni 1000 flood vtep 30.30.30.1
```

SW2 Configure

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
set vlan-interface loopback address 20.20.20.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.2.1 prefix-length 24
set protocols static route 10.10.10.1/32 next-hop 100.10.2.2
set vxlans source-interface loopback address 20.20.20.1
set vxlans vni 1000 interface te-1/1/10 vlan 10
set vxlans vni 1000 flood vtep 10.10.10.1
```

SW3 Configure

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
set vlan-interface loopback address 30.30.30.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.3.1 prefix-length 24
set protocols static route 10.10.10.1/32 next-hop 100.10.3.2
set vxlans source-interface loopback address 30.30.30.1
set vxlans vni 1000 interface te-1/1/10 vlan 10
set vxlans vni 1000 flood vtep 10.10.10.1
```

VXLAN ECMP Configuration

Requirements

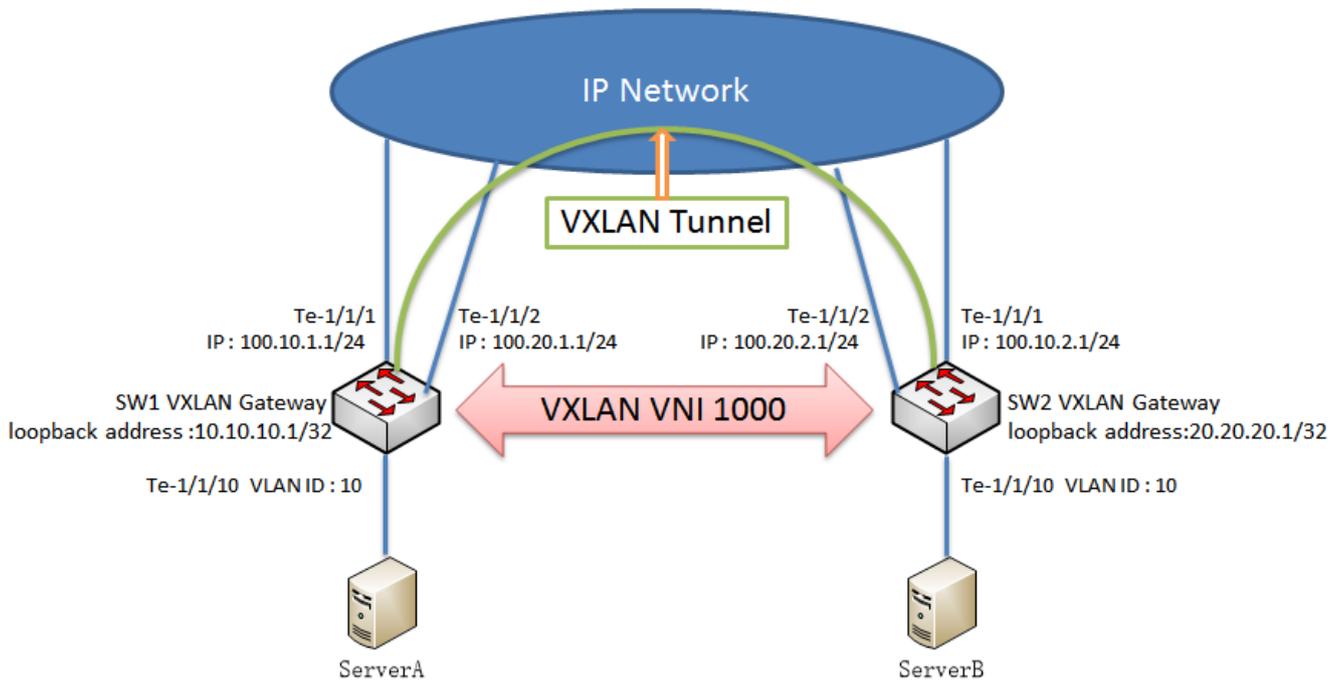
This example uses the following hardware and software components:

- An Trident-II Series switch
- PICOS OS Release more than 2.5 version

Overview

In this example, VXLAN ECMP is configured to run on a VXLAN domain. VTEP interfaces sources are configured to the loopback address. Interfaces are configured for VLAN tagging and encapsulation. Static route protocols is configured to facilitate unicast routing.

Topology



The VXLAN that comprise the networks include:

- SW1 VNI 1000: VTEP 10.10.10.1: VLAN 10
- SW2 VNI 1000: VTEP 20.20.20.1: VLAN 10
- SW1 IP Adress: 100.10.1.1 and IP Address: 100.20.1.1
- SW2 IP Adress: 100.10.2.1 and IP Address: 100.20.2.1

Configuring VXLAN on Trident-II Series Switches

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and copy and paste the commands into the CLI.

SW1 Configure

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set vlans vlan-id 2000 l3-interface 2000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/2 family ethernet-switching native-vlan-id 2000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
set vlan-interface loopback address 10.10.10.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.1.1 prefix-length 24
set vlan-interface interface 2000 vif 2000 address 100.20.1.1 prefix-length 24
set protocols static route 20.20.20.1/32 next-hop 100.10.1.2
set protocols static route 20.20.20.1/32 qualified-next-hop 100.20.1.2 metric 1
```

```
set vxlans source-interface loopback address 10.10.10.1
set vxlans vni 1000 interface te-1/1/10 vlan 10
set vxlans vni 1000 flood vtep 20.20.20.1
```

SW2 Configure

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set vlans vlan-id 2000 l3-interface 2000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/2 family ethernet-switching native-vlan-id 2000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
set vlan-interface loopback address 20.20.20.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.2.1 prefix-length 24
set vlan-interface interface 2000 vif 2000 address 100.20.2.1 prefix-length 24
set protocols static route 10.10.10.1/32 next-hop 100.10.2.2
set protocols static route 10.10.10.1/32 qualified-next-hop 100.20.2.2 metric 1
set vxlans source-interface loopback address 20.20.20.1
set vxlans vni 1000 interface te-1/1/10 vlan 10
set vxlans vni 1000 flood vtep 10.10.10.1
```

Configuring VXLAN Step-by-Step Procedure

The following example shows how to set up a basic VXLAN ECMP configuration with VXLAN domain. To configure VXLAN ECMP on Trident-II Series switches, follow these steps:

Configure VXLAN Step-by-Step for SW1

1. Configure the VLAN ID to 10 for vxlan domain.

```
set vlans vlan-id 10
```

2. Configure the VLAN ID to 1000 and 2000 for ip routing.

```
set vlans vlan-id 1000 l3-interface 1000
set vlans vlan-id 2000 l3-interface 2000
```

3. Configure the te-1/1/1 interface VLAN ID to 1000.

```
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
```

4. Configure the te-1/1/2 interface VLAN ID to 2000.

```
set interface gigabit-ethernet te-1/1/2 family ethernet-switching native-vlan-id 2000
```

5. Configure the te-1/1/10 interface VLAN ID to 10

```
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
```

6. Configure ip address for the loopback interface.

```
set vlan-interface loopback address 10.10.10.1 prefix-length 32
```

7. Configure ip address for the vlan-interface 1000.

```
set vlan-interface interface 1000 vif 1000 address 100.10.1.1 prefix-length 24
```

8. Configure ip address for the vlan-interface 2000.

```
set vlan-interface interface 2000 vif 2000 address 100.20.1.1 prefix-length 24
```

9. Configure static route for the VXLAN ECMP.

```
set protocols static route 20.20.20.1/32 next-hop 100.10.1.2
set protocols static route 20.20.20.1/32 qualified-next-hop 100.20.1.2 metric 1
```

10. Configure VTEP interface sources ip address.

```
set vxlans source-interface loopback address 10.10.10.1
```

11. Configure a VLAN ID 10 to a VNI 1000

```
set vxlans vni 1000 interface te-1/1/10 vlan 10
```

12. Configure flood vtep ip address for a VNI

```
set vxlans vni 1000 flood vtep 20.20.20.1
```

Configure VXLAN Step-by-Step for SW2

1. Configure the VLAN ID to 10 for vxlan domain.

```
set vlans vlan-id 10
```

2. Configure the VLAN ID to 1000 and 2000 for ip routing.

```
set vlans vlan-id 1000 13-interface 1000
set vlans vlan-id 2000 13-interface 2000
```

3. Configure the te-1/1/1 interface VLAN ID to 1000.

```
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
```

4. Configure the te-1/1/2 interface VLAN ID to 2000.

```
set interface gigabit-ethernet te-1/1/2 family ethernet-switching native-vlan-id 2000
```

5. Configure the te-1/1/10 interface VLAN ID to 10

```
set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 10
```

6. Configure ip address for the loopback interface.

```
set vlan-interface loopback address 20.20.20.1 prefix-length 32
```

7. Configure ip address for the vlan-interface 1000.

```
set vlan-interface interface 1000 vif 1000 address 100.10.2.1 prefix-length 24
```

8. Configure ip address for the vlan-interface 2000.

```
set vlan-interface interface 2000 vif 2000 address 100.20.2.1 prefix-length 24
```

9. Configure static route for the VXLAN ECMP.

```
set protocols static route 10.10.10.1/32 next-hop 100.10.2.2
set protocols static route 10.10.10.1/32 qualified-next-hop 100.20.2.2 metric 1
```

10. Configure VTEP interface sources ip address.

```
set vxlans source-interface loopback address 20.20.20.1
```

11. Configure a VLAN ID 10 to a VNI 1000

```
set vxlans vni 1000 interface te-1/1/10 vlan 10
```

12. Configure flood vtep ip address for a VNI

```
set vxlans vni 1000 flood vtep 10.10.10.1
```

View the VXLAN Table of SW1

```
admin@XorPlus# run show vxlan
Egress map:
egress_id 100006 MAC 22:22:22:22:22:1, port_id 1/1/1, vif_index 7 unicast
egress_id 100008 MAC 22:22:22:22:22:2, port_id 1/1/2, vif_index 8 unicast
L3 tunnel mac map:
vlan id 1000, ref_count 1
vlan id 2000, ref_count 1
Port vlan map mode map:
port id 1/1/10, ref_count 1
Termination admin state map:
port id 1/1/10, ref_count 1
Tunnel Map:
```

```
tunnel id 0x4C000001, dst_vtep 20.20.20.1, nexthops (100.10.1.2 100.20.1.2 ), ecmp_id
200256, ref_count 1
tunnel id 0, dst_vtep 224.0.0.1, nexthops (), ecmp_id 0, ref_count 1
Access ports:
id 0x80000001, vpn_id 0x7000, port_id 1/1/10, vlan_id 10, egress_id 100005
Network ports:
id 0x80000002, vpn_id 0x7000, port_id 1/1/1, egress_id 100006, tunnel_id 0x4C000001,
unicast
id 0x80000003, vpn_id 0x7000, port_id 1/1/1, egress_id 100007, tunnel_id 0x4C000001,
multicast
BFD sessions:
admin@XorPlus#
```

The `show vxlan` displays information about VXLAN endpoint configuration. Display VXLAN ECMP to the remote vxlan tunnel vtep has two nexthops(100.10.1.2 100.20.1.2).

OVSDB VTEP Configuration

- Configuring an OVSDB VTEP
- OVSDB VTEP with Midonet Configuration
- OVSDB VTEP with NSX Configuration
- OVSDB VTEP with vtep-ctl Configuration Examples

Configuring an OVSDB VTEP

Overview

A Pica8 Network device implements Open vSwitch Database (OVSDB) management protocol, has connections with network virtualization controller, and can use Open vSwitch Database management protocol to deploy and manage VXLANs on VTEPs. For more OVSDB management protocol details, please read RFC7047.

Configuring an OVSDB Connection with a Controller

The device supports the following types OVSDB connections:

- SSL —The device initiates an SSL connection to a network virtualization controller.
- PSSSL —The device accepts the SSL connection from a network virtualization controller.
- TCP —The device initiates a TCP connection to a network virtualization controller.
- PTCP —The device accepts the TCP connection from a network virtualization controller.

Creating and Installing an SSL Key and Certificate on a Pica8 Device

To secure a connection between a Pica8 device that supports the Open vSwitch Database (OVSDB) management protocol and one or more network virtualization controllers, the following Secure Sockets Layer (SSL) files must be present in the `/ovs/var/lib/openvswitch/pki` directory on the device:

- `pica8-cert.pem`
- `pica8-privkey.pem`

- pica8-req.pem

You must create the pica8-privkey.pem and pica8-cert.pem files for the device, and then install the two files in the /ovs/var/lib/openvswitch/pki directory on the device.

To create and install an SSL key and certificate on a Pica device:

1. Update the system time, or certificate will show expired.
2. Install openssl on Pica8 device or linux computer
3. Initialize a PKI on Pica8 device or linux computer
ovs-pki init --force
4. On Pica8 device or the same Linux computer on which the PKI exists, create a new key and certificate for the Pica8 device.
ovs-pki req+sign pica8
5. Copy only the pica8-privkey.pem and pica8-cert.pem files from the Linux computer to the /ovs/var/lib/openvswitch/pki directory on the Pica8 device.

Configure a OVSDb VTEP Step

OVSDb VTEP supported on PicOS L2/L3 switch, the commands configure step as below.

Enabling the OVSDb server

```
set vxlans ovbdb-managed true
```

Configuring a source address for VXLAN tunnels

```
set vxlans source-interface loopback address 10.10.10.1
```

Disabling VXLAN tunnels address learning

```
set vxlans tunnel-mac-learning disable true
```

Configuring the VTEP management interface IP address

```
set protocols ovbdb management-ip 10.10.51.157
```

Configuring an SSL connection to a network virtualization controller

```
set protocols ovbdb SSL private-key "/ovs/var/lib/openvswitch/pki/pica8-privkey.pem"
set protocols ovbdb SSL certificate "/ovs/var/lib/openvswitch/pki/pica8-cert.pem"
set protocols ovbdb SSL ca-cert "/ovs/var/lib/openvswitch/pki/controller.cacert"
set protocols ovbdb SSL bootstrap true
set protocols ovbdb controller c1 address 10.10.50.220
set protocols ovbdb controller c1 protocol ssl
```

Configuring PSSSL connection requests from a network virtualization controller

```
set protocols ovssdb SSL private-key "/ovs/var/lib/openvswitch/pki/pica8-privkey.pem"  
set protocols ovssdb SSL certificate "/ovs/var/lib/openvswitch/pki/pica8-cert.pem"  
set protocols ovssdb SSL ca-cert "/ovs/var/lib/openvswitch/pki/controller.cacert"  
set protocols ovssdb SSL bootstrap true  
set protocols ovssdb controller c1 protocol ssl
```

Configuring TCP connection to a network virtualization controller

```
set protocols ovssdb controller c1 address 10.10.50.220  
set protocols ovssdb controller c1 protocol tcp
```

Configuring PTCP connection requests from a network virtualization controller

```
set protocols ovssdb controller c1 protocol ptcp
```

OVSSDB VTEP with Midonet Configuration

Requirements

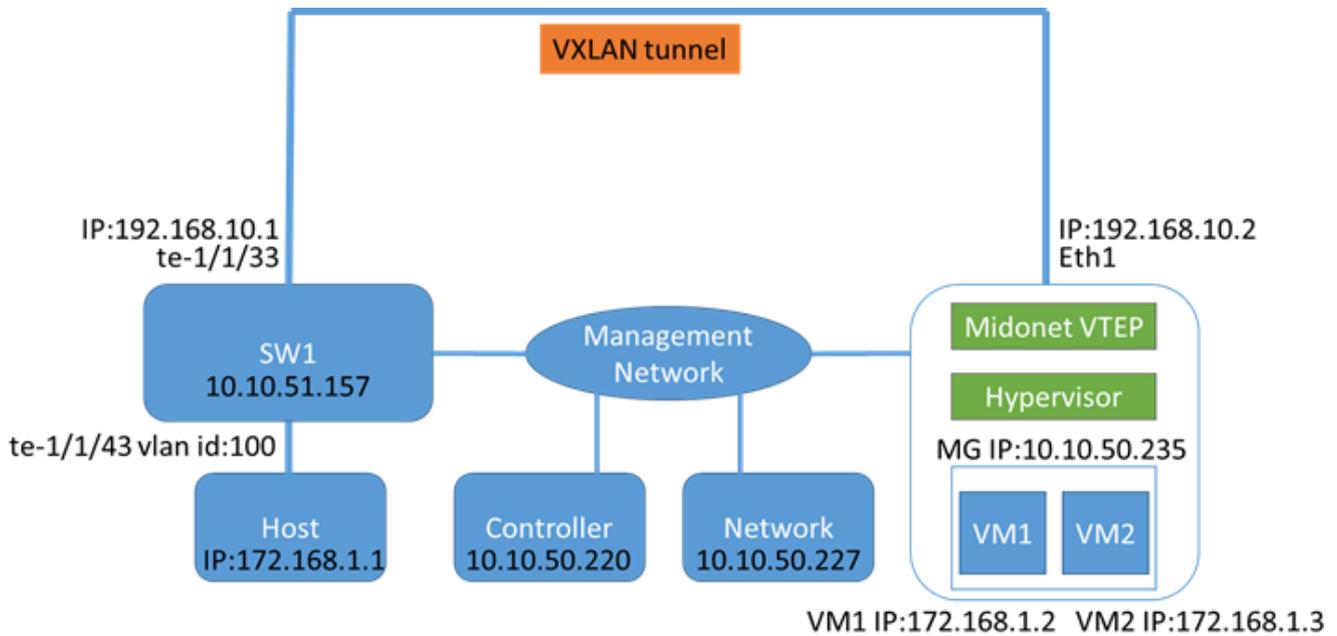
This example uses the following hardware and software components:

- Trident-II Series switch
- PicOS OS Release 2.6 and higher
- Midonet controller

Overview

In this example, OVSSDB VTEP is configured to run on a VXLAN domain. VTEP interface sources are configured to the loopback address. Interfaces are configured for VLAN tagging and encapsulation. Static route protocols are configured to facilitate unicast routing.

Topology



The VXLAN networks Host and VM IP address and Mac address information:

Host IP Address: 172.168.1.1, Mac address: 00:07:43:05:45:11

VM1 IP Address: 172.168.1.2, Mac address: fa:16:3e:00:0c:f3

VM2 IP Address: 172.168.1.3, Mac address: fa:16:3e:28:aa:cd

Configuring OVSDb on Trident-II Series Switches

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and copy and paste the commands into the CLI.

SW1 Configure :

```
set interface ge-interface-mode "SFP"
set interface gigabit-ethernet te-1/1/33 speed "1000"
set interface gigabit-ethernet te-1/1/33 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/43 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet te-1/1/43 family ethernet-switching vlan members 100
set protocols ovssdb management-ip 10.10.51.157
set protocols ovssdb controller cl protocol "ptcp"
set protocols ovssdb interface te-1/1/43
set vlan-interface interface 1000 vif 1000 address 192.168.10.1 prefix-length 24
set vlans vlan-id 100
set vlans vlan-id 1000 l3-interface "1000"
set vxlans source-interface 1000 address 192.168.10.1
set vxlans ovssdb-managed true
```

Configuring OVSDB Step-by-Step Procedure

The following example shows how to set up a basic VTEP OVSDB configuration with VXLAN domain. To configure VXLAN on Trident-II Series switches, follow these steps:

Configuring OVSDB Step-by-Step for SW1

1. Configure the VLAN ID to 100 for vxlan domain.

```
set vlans vlan-id 100
```

2. Configure the VLAN ID to 1000 for IP routing.

```
set vlans vlan-id 1000 13-interface vlan-1000
```

3. Configure the te-1/1/1 interface VLAN ID to 1000.

```
set interface gigabit-ethernet te-1/1/33 family ethernet-switching native-vlan-id 1000
```

4. Configure VLAN trunk for te-1/1/43.

```
set interface gigabit-ethernet te-1/1/43 family ethernet-switching port-mode "trunk"
```

5. Configure the te-1/1/43 interface VLAN ID to 100.

```
set interface gigabit-ethernet te-1/1/43 family ethernet-switching vlan members 100
```

6. Configure IP address for the loopback interface.

```
set vlan-interface loopback address 10.10.10.1 prefix-length 32
```

7. Configure IP address for the vlan-interface vlan-1000.

```
set vlan-interface interface vlan-1000 vif vlan-1000 address 192.168.10.1 prefix-length 24
```

8. Configure VTEP interface sources IP address.

```
set vxlans source-interface vlan-1000 address 192.168.10.1
```

9. Enable VXLAN managed by ovsdb

```
set vxlan ovsdb-managed true
```

10. Configure ovssdb management interface IP address

```
set protocols ovssdb management-ip 10.10.51.157
```

11. Configure ovssdb controller protocol

```
set protocols ovssdb controller c1 protocol ptcp
```

12. Configure ovssdb controller port

```
set protocols ovssdb controller ovssdb port 6632
```

13. Configure ovssdb the interface on the VTEP

```
set protocols ovssdb interface te-1/1/43
```

Manually Configuring the Switch as a VTEP Gateway by Midonet cli

Configuring the ovssdb by midonet cli for SW1

1. Add a virtualization image on OpenStack Dashboard

Images



Project (1) Shared with Me (0) Public (0) + Create Image X Delete Images

Image Name	Type	Status	Public	Protected	Format	Size	Actions
test	Image	Active	No	No	QCOW2	12.6 MB	Launch

Displaying 1 item

2. Create a network on OpenStack Dashboard

Networks



+ Create Network X Delete Networks

Project	Network Name	Subnets Associated	DHCP Agents	Shared	Status	Admin State	Actions
admin	vxlan	n1 172.168.1.0/24	1	No	ACTIVE	UP	Edit Network

Displaying 1 item

3. Create two virtual hosts, and add them to network on OpenStack Dashboard

Instances

Instances

Instance Name Filter Filter Launch Instance Soft Reboot Instances Terminate Instances

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
vm2	test	172.168.1.3	m1.tiny	-	Active	nova	None	Running	3 hours, 19 minutes	Create Snapshot
vm1	test	172.168.1.2	m1.tiny	-	Active	nova	None	Running	3 hours, 21 minutes	Create Snapshot

Displaying 2 items

4. Create a tunnel zone of type 'vtep'

```
midonet> tunnel-zone create name vtep_zone1 type vtep
tzone0
```

5. Add a VTEP to MidoNet, and assign it to the 'vtep' tunnel zone that you created.

```
midonet> vtep add management-ip 10.10.51.157 management-port 6632 tunnel-zone tzone0
name br0 description OVS VTEP Emulator management-ip 10.10.51.157 management-port
6632 tunnel-zone tzone0 connection-state CONNECTED
midonet> list vtep
name br0 description management-ip 10.10.51.157 management-port 6632 tunnel-zone
tzone0 connection-state CONNECTED
```

6. Create a binding between the VTEP and a Neutron network behind a MidoNet bridge

```
midonet> bridge list
bridge br1 name vxlan state up
midonet> show bridge br1 id
85296f07-2235-4963-8160-fb66eca85675
midonet>
```

7. Add the host's IP address to the same tunnel zone as the VTEP

```
midonet> tunnel-zone tzone0 add member host host0 address 192.168.10.2
zone tzone0 host host0 address 192.168.10.2
midonet>
```

8. Create a binding between the VTEP's vlan 100 interface te-1/1/43 and the Neutron network behind the bridge1

```
midonet> vtep management-ip 10.10.51.157 binding add network-id
85296f07-2235-4963-8160-fb66eca85675 physical-port te-1/1/43 vlan 100
Internal error: The server could not comply with the request since it is either
malformed or otherwise incorrect.
midonet> vtep management-ip 10.10.51.157 binding list
management-ip 10.10.51.157 physical-port te-1/1/43 vlan 100 network-id
85296f07-2235-4963-8160-fb66eca85675
```

9. Add the IP address of the host on the VTEP to the security group ip-address-group0

```
midonet>
```

```

midonet> ip-address-group ip-address-group0 add ip address 172.168.1.1
address 172.168.1.1
midonet>

```

View the VXLAN table of SW1:

```

admin@XorPlus# run show vxlan
Egress map:
    egress_id 100009 MAC 0:c:29:23:31:9, port_id 1/1/33, vif_index 8 unicast
L3 tunnel mac map:
    vlan id 1000, ref_count 1
Port vlan map mode map & Termination admin state map:
    port id 1/1/43, ref_count 1
Tunnel Map:
    tunnel id 0X4C000200, dst_vtep 192.168.10.2, nexthops (192.168.10.2 ), ecmp_id
100009, ref_count 1
    tunnel id 0X4C000001, dst_vtep 224.0.0.1, nexthops (), ecmp_id 0, ref_count 1
Access ports:
    id 0X80000002, vpn_id 0X7000, port_id 1/1/43, vlan_id 100, egress_id 100010
Network ports:
    id 0X80000003, vpn_id 0X7000, port_id 1/1/33, egress_id 100009, tunnel_id
0X4C000200, unicast
    id 0X80000004, vpn_id 0X7000, port_id 1/1/33, egress_id 100011, tunnel_id
0X4C000200, multicast
    id 0X80000001, vpn_id 0XFFFFFFFF, port_id 1/1/0, egress_id 100006, tunnel_id
0X4C000001, multicast
BFD sessions:
admin@XorPlus#

```

The show vxlan to displays information about VXLAN endpoint configuration. Display VXLAN to the remote vxlan tunnel vtep has a nexthops(192.168.10.2).

View the VXLAN mac table of SW1:

```

admin@XorPlus# run show vxlan address-table
VNID          MAC address          Type          Interface          VTEP
-----
10001         00:1e:c9:bb:bb:ce    Dynamic      te-1/1/43
10001         fa:16:3e:00:0c:f3    Static       192.168.10.2
10001         fa:16:3e:28:aa:cd    Static       192.168.10.2
admin@XorPlus#

```

Dump the ovssdb hardware vtep table of SW1:

```

root@XorPlus$ovsdb-client dump hardware_vtep
Arp_Sources_Local table
_uuid locator src_mac
-----

Arp_Sources_Remote table
_uuid locator src_mac
-----

Global table
_uuid          managers          switches
-----
4146166b-ad2e-4d05-857f-8ba4b3f0ac0d [bd6ac790-b304-4ed7-a77b-8ab7063b8132]

```

```
[cfdcc9fa-0295-44b0-81c3-c975b3d463cb]

Logical_Binding_Stats table
_uuid bytes_from_local bytes_to_local packets_from_local packets_to_local
-----

Logical_Router table
_uuid description name static_routes switch_binding
-----

Logical_Switch table
_uuid description name
options tunnel_key
-----
-----
ele37b4a-37fe-43f2-a9f7-3a9925b6e92e "" "mn-85296f07-2235-4963-8160-fb66eca85675"
{} 10001

Manager table
_uuid inactivity_probe is_connected max_backoff other_config
status target
-----
-----
bd6ac790-b304-4ed7-a77b-8ab7063b8132 30000 true 3000 {}
{bound_port="6632", sec_since_connect="13921", state=ACTIVE} "ptcp:6632"

Mcast_Macs_Local table
MAC _uuid ipaddr locator_set logical_switch
-----

Mcast_Macs_Remote table
MAC _uuid ipaddr locator_set
logical_switch
-----
-----
unknown-dst 8c6b4993-7be8-4d85-811b-3255e15d2f92 ""
6d275247-2c1b-4c79-8f08-b17d93bd1e32 ele37b4a-37fe-43f2-a9f7-3a9925b6e92e

Physical_Locator table
_uuid dst_ip encapsulation_type
-----
-----
d983943f-c791-4431-89a2-ec6a531a4d15 "192.168.10.1" "vxlan_over_ipv4"
09c0f3c2-d42a-406b-8644-3bffc472a247 "192.168.10.2" "vxlan_over_ipv4"

Physical_Locator_Set table
_uuid locators
-----
-----
6d275247-2c1b-4c79-8f08-b17d93bd1e32 [09c0f3c2-d42a-406b-8644-3bffc472a247]

Physical_Port table
_uuid description name port_fault_status
vlan_bindings vlan_stats
-----
-----
35f008a2-e248-4330-a1e5-85f3f843bc68 "" "te-1/1/43" []
{100=ele37b4a-37fe-43f2-a9f7-3a9925b6e92e} {}

Physical_Switch table
_uuid description management_ips name ports
switch_fault_status tunnel_ips tunnels
-----
-----
-----
cfdcc9fa-0295-44b0-81c3-c975b3d463cb "" ["10.10.51.157"] "br0"
[35f008a2-e248-4330-a1e5-85f3f843bc68, 3d5eae61-46bc-4e3c-84f3-06aed7961ff5,
c30d1ef4-b54a-4946-bd8c-460af234875e] [] ["192.168.10.1"] []

SSL table
_uuid bootstrap_ca_cert ca_cert certificate external_ids private_key
-----
```

```

Tunnel table
 _uuid bfd_config_local bfd_config_remote bfd_params bfd_status local remote
-----

Ucast_Macs_Local table
MAC          _uuid          ipaddr locator
logical_switch
-----

"00:1e:c9:bb:bb:ce" f431e446-6c1c-4842-ad3a-19cd04a54952 ""
d983943f-c791-4431-89a2-ec6a531a4d15 e1e37b4a-37fe-43f2-a9f7-3a9925b6e92e

Ucast_Macs_Remote table
MAC          _uuid          ipaddr locator
logical_switch
-----

"fa:16:3e:00:0c:f3" 0c6732d3-9e72-4444-9d8f-07abde993aa7 ""
09c0f3c2-d42a-406b-8644-3bffc472a247 e1e37b4a-37fe-43f2-a9f7-3a9925b6e92e
"fa:16:3e:28:aa:cd" b1160e14-8e35-4293-a987-59ddc29f7304 ""
09c0f3c2-d42a-406b-8644-3bffc472a247 e1e37b4a-37fe-43f2-a9f7-3a9925b6e92e
root@XorPlus$

```

Ping VM1 and VM2 on the Host:

```

root@Dev-45:~# ping 172.168.1.2 -c 5
PING 172.168.1.2 (172.168.1.2) 56(84) bytes of data.
64 bytes from 172.168.1.2: icmp_req=1 ttl=64 time=3.92 ms
64 bytes from 172.168.1.2: icmp_req=2 ttl=64 time=1.51 ms
64 bytes from 172.168.1.2: icmp_req=3 ttl=64 time=1.47 ms
64 bytes from 172.168.1.2: icmp_req=4 ttl=64 time=1.59 ms
64 bytes from 172.168.1.2: icmp_req=5 ttl=64 time=1.57 ms

--- 172.168.1.2 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 1.476/2.015/3.920/0.954 ms
root@Dev-45:~# ping 172.168.1.3 -c 5
PING 172.168.1.3 (172.168.1.3) 56(84) bytes of data.
64 bytes from 172.168.1.3: icmp_req=1 ttl=64 time=10.1 ms
64 bytes from 172.168.1.3: icmp_req=2 ttl=64 time=1.70 ms
64 bytes from 172.168.1.3: icmp_req=3 ttl=64 time=1.64 ms
64 bytes from 172.168.1.3: icmp_req=4 ttl=64 time=1.62 ms
64 bytes from 172.168.1.3: icmp_req=5 ttl=64 time=1.67 ms

--- 172.168.1.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 1.623/3.365/10.189/3.412 ms
root@Dev-45:~# arp -n

```

Address	HWtype	HWaddress	Flags Mask	Iface
172.168.1.3	ether	fa:16:3e:28:aa:cd	C	eth1.100
172.168.1.2	ether	fa:16:3e:00:0c:f3	C	eth1.100

OVSDB VTEP with NSX Configuration

Requirements

This example uses the following hardware and software components:

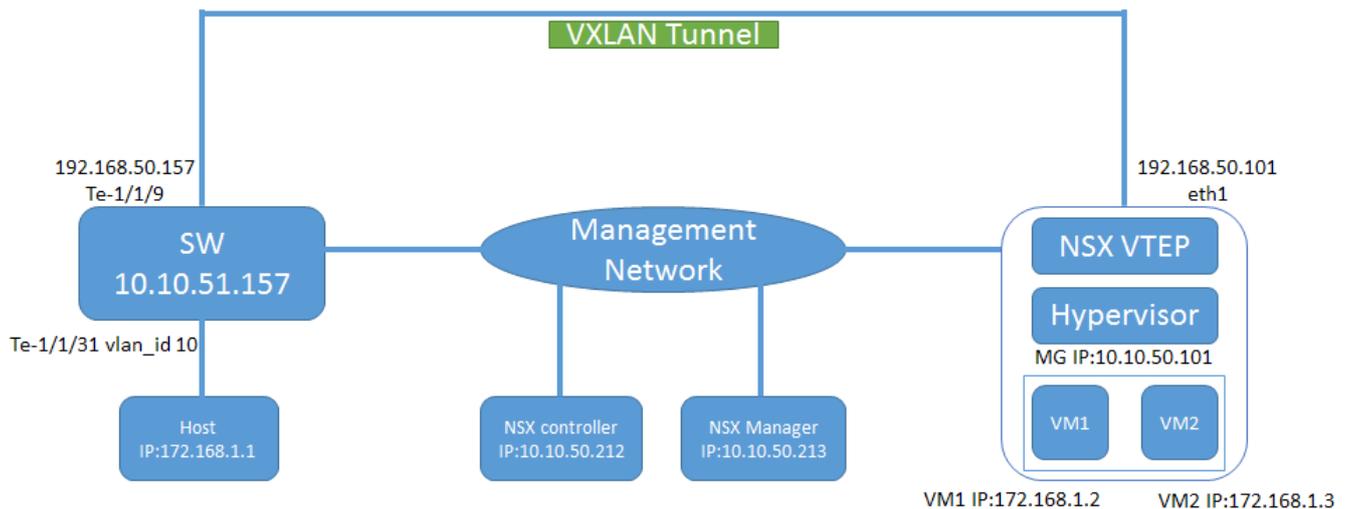
- A Trident-II Series switch
- PicOS (more than 2.6 version)

- NSX controller

Overview

In this example, OVSDb VTEP is configured to run on a VXLAN domain. VTEP interfaces sources are configured to the loopback address. Interfaces are configured for VLAN tagging and encapsulation. Static route protocols are configured to facilitate unicast routing.

Topology



The VXLAN networks Host and VM IP address and Mac address information :

Host IP Address: 172.168.1.1, Mac address: 70:72:cf:9d:6f:fb

VM1 IP Address: 172.168.1.2, Mac address: 00:50:56:ae:46:d3

VM2 IP Address: 172.168.1.3, Mac address: 00:50:56:ae:5a:30

Configuring OVSDb on Trident-II Series Switches

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and copy and paste the commands into the CLI.

SW1 Configure

```
set interface gigabit-ethernet te-1/1/9 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/31 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet te-1/1/31 family ethernet-switching vlan members 100
set protocols ovssdb management-ip 10.10.51.157
set protocols ovssdb controller ovssdb protocol ssl
set protocols ovssdb controller ovssdb address 10.10.50.212
set protocols ovssdb controller ovssdb port 6632
set protocols ovssdb interface te-1/1/31
set vlan-interface interface 1000 vif 1000 address 192.168.50.157 prefix-length 24
set vlans vlan-id 100
```

```
set vlans vlan-id 1000 l3-interface "1000"
set vxlans source-interface 1000 address 192.168.50.157
set vxlans ovssdb-managed true
```

Configuring OVSSDB Step-by-Step Procedure

The following example shows how to set up a basic VTEP OVSSDB configuration with VXLAN domain. To configure VXLAN on an Trident-II Series switches, follow these steps:

Configure OVSSDB Step-by-Step for SW1

1. Configure the VLAN ID to 100 for vxlan domain.

```
set vlans vlan-id 100
```

2. Configure the VLAN ID to 1000 for ip routing.

```
set vlans vlan-id 1000 l3-interface vlan-1000
```

3. Configure the te-1/1/9 interface VLAN ID to 1000.

```
set interface gigabit-ethernet te-1/1/9 family ethernet-switching native-vlan-id 1000
```

4. Configure VLAN trunk for te-1/1/31.

```
set interface gigabit-ethernet te-1/1/31 family ethernet-switching port-mode "trunk"
```

5. Configure the te-1/1/31 interface VLAN ID to 100.

```
set interface gigabit-ethernet te-1/1/31 family ethernet-switching vlan members 100
```

6. Configure IP address for the vlan-interface vlan-1000.

```
set vlan-interface interface vlan-1000 vif vlan-1000 address 192.168.50.157
prefix-length 24
```

7. Configure VTEP interface sources IP address.

```
set vxlans source-interface vlan-1000 address 192.168.50.157
```

8. Enable VXLAN managed by ovssdb

```
set vxlan ovssdb-managed true
```

9. Configure ovssdb controller ip address

```
set protocols ovssdb controller ovssdb address 10.10.50.212
```

10. Configure ovssdb controller protocol

```
set protocols ovssdb controller ovssdb protocol ssl
```

11. Configure ovssdb controller port

```
set protocols ovssdb controller ovssdb port 6632
```

12. Configure ovssdb the interface on the VTEP

```
set protocols ovssdb interface te-1/1/31
```

Manually configuring the Switch as a VTEP Gateway by NSX Manager

Configuring the ovssdb by NSX Manager for SW1

1. Add components in NSX manager.

The screenshot displays the VMware NSX Manager interface. The top navigation bar includes 'Dashboard', 'Network Components', 'Controller Cluster', and 'Tools & Troubleshooting'. The user is logged in as 'Welcome, admin' and the current cluster is 'NSX-Controller-...'. The dashboard provides a comprehensive overview of the system's health and configuration.

Summary of Logical Components

Type	Registered	Active	Alert	Action
Switches	1	1	0	Add
Switch Ports	5	4	0	Add
Routers	0	0	0	Add
Router Ports	0	0	0	Add
Queues	0	0	0	Add
ACLs	0	0	0	Add
Security Profiles	0	0	0	Add

Summary of Transport Components

Type	Registered	Active	Alert	Action
Gateways	1	0	1	Add
Hypervisors	2	2	0	Add
Service Nodes	0	0	0	Add
Zones	1	1	0	Add
Gateway Services	1	N/A	N/A	Add

Hypervisor Software Version Summary

Platform Type	Platform Version	vSwitch Version	Count
ESXi	5.5.0-Release-build-1331820	2.3.0.38094	2

NSX Appliance Software Version Summary

Appliance Type	NSX Version	Count
- Unknown -	- Unknown -	1
NSX Controller Node	4.0.1.30244	1

Controller Cluster Nodes

Name	Management Address	Status	System Load	System Memory Usage	File System Usage
50e87c0a-d677-44b8-846c-621f0d3429af	10.10.50.213	Up	7%	37%	67% /boot 16%

Service Nodes

No Service Nodes

Multi-Domain Interconnect Summary

No Multi-Domain Interconnect Gateway Services

2. Add ESXi in NSX Manager

Transport Node "vxlan-101" [Edit](#) [Delete](#)

Hypervisor

Hypervisor
🌐 ↻ ⏸

Properties

UUID: c2059471-30b5-4c3c-ba5e-073f81bda51f

Tags: - None -

System Type: ESXi

System Version: 5.5.0-Release-build-1331820

OVS Version: 2.3.0.38094

Integration Bridge: br-int ● Found

Tunnel Keep-alive Spray: ● No

Credential: 10.10.50.101

Master Cluster Node: [50e87c0a-d677-44b8-846c-621f0d3429af](#)

Status

Management Connection: ● Up (10.10.50.101)

OpenFlow Connection: ● Up

Admin Status: ✔ Enabled

System Statistics

System Memory (MB) (used/cache/total): 0 / 0 / 0

System Swap (MB) (used/total): 0 / 0

File System Usage (MB) (path:used/total): N/A

System Uptime: -

CPU Cores: 0

System Load Average: -

Last Connected: Feb 4, 2015 2:33 GMT

Last Disconnected: Feb 4, 2015 2:33 GMT

Process Statistics

Restart Count: -

OVS Uptime: -

OVS Monitor Uptime: -

OVS CPU Used: -

Transport Connectors

Transport Zone	Transport Zone UUID	Type	Bridge ID	IP Address
vxlan-10010	f1f07...d5152	VXLANConnector	-	192.168.50.101

3. Add hardware gateway

vmware
NSX

Dashboard
Network Components
Controller Cluster
Tools & Troubleshooting

Network Components > Gateway Services

Gateway Service "vxlan-11" [Edit](#) [Delete](#)

Gateway Service
🌐 ↻ ⏸

Properties

UUID: 1fdcc6ee-2382-4a56-9c65-8f0a61d808ef

Tags: - None -

Type: VTEP L2

Gateways

Transport Node UUID	Transport Node Name	Port ID
8a4d2284-e495-4ab1-9c71-83df74f1d8a4	vxlan-11	te-1/1/31
8a4d2284-e495-4ab1-9c71-83df74f1d8a4	vxlan-11	te-1/1/32
8a4d2284-e495-4ab1-9c71-83df74f1d8a4	vxlan-11	ae1

▶ Recent Logs (12)

▶ Logical Switch Ports (0)

594

4. Create Transport Zone

Network Components > Transport Zones

Transport Zone "vxlan-10010" Edit Delete

Transport Zone

Properties

UUID f1f074e3-e14c-49b8-9137-e098b00d5152

Tags - None -

Transport Node Mo... Options

5. Create Logical switch

Logical Switch "vxlan1" Edit Delete Port Connections

Logical Switch

Properties

UUID 4677abd6-84c1-4aa1-b27b-c06c15eb4b58

Tags - None -

Port Isolation Disabled

Replication Mode Source Node

Transport Zone Bindings

Transport Zone UUID	Transport Zone Name	Type	Config
f1f07...d5152	vxlan-10010	VXLAN	VNI: 10010

Status & Statistics

Fabric Status Up

Configured Ports 5

Admin Status Up/Down 5 / 0

Link Status Up/Down 4 / 1

Logical Router

No attached logical router

Change

Multi-Domain Interconnect Gateway Services

Name	UUID	Context ID
No Multi-Domain Interconnect Gateway Services		

Transport Node Mo... Options

6. Create Logical port

Logical Switch Port (5)

Filter Delete Checked Add

	Admin	Link	Fabric	Message	Name	Port	Switch Name	UUID	Attachment	Attached MAC
<input type="checkbox"/>	Up	Up	Up	-	vxlan-203	1	vxlan1	4edb8...b203d	VIF: 502e89d0-baec-a0c8-049a-55b6b1da4099-0	00:50:56:ae:46:d3
<input type="checkbox"/>	Up	Up	Up	-	vxlan-201	2	vxlan1	fd141...af6ba	VIF: 502e6aa0-311c-5a42-48f7-4e2253cd5fb7-0	00:50:56:ae:f8:73
<input type="checkbox"/>	Up	Up	Up	-	vxlan-204	4	vxlan1	c6494...888fa	VIF: 502ea58b-34f6-159a-ca8a-d5d6fdcef7ad-0	00:50:56:ae:5a:30
<input type="checkbox"/>	Up	Up	Up	-	vxlan-202	5	vxlan1	d146b...9706b	VIF: 502e2430-ce49-5ccb-3c4c-04aa292f273d-0	00:50:56:ae:d9:db

7. Add gateway VTEP port

Logical Switch Port "6a275791-299b-4893-9ac0-ec0ad3511522" Edit Delete

Logical Switch "vxlan1"

View the VXLAN table of SW1:

```
admin@XorPlus# run show vxlan tunnel
Total number of tunnels: 1
VNI 2, Encap:service-vlan-delete, Decap:service-vlan-add
src addr:192.168.50.157, dst addr:192.168.50.101, state:UP
traffic type:all
nexthops:192.168.50.101
output ports:te-1/1/9
```

Show vxlan to display information about VXLAN endpoint configuration. Display VXLAN to the remote vxlan tunnel vtep has a nexthops(192.168.50.101).

View the VXLAN mac table of SW1:

```
admin@XorPlus# run show vxlan address-table
VNID          MAC address          Type          Interface          VTEP
-----
10010         70:72:cf:9d:6f:fb    Dynamic       te-1/1/31
10010         00:50:56:ae:46:d3    Static        192.168.50.101
10010         00:50:56:ae:5a:30    Static        192.168.50.101
admin@XorPlus#
```

Dump the ovssdb hardware vtep table of SW1:

```
root@XorPlus$ovsdb-client dump hardware_vtepArp_Sources_Local table
_uuid locator src_mac
-----
Arp_Sources_Remote table
_uuid locator src_mac
-----
Global table
_uuid managers switches
-----
-----
3542b066-e49a-4df6-91e5-731ce43868c7 [5cc533b6-3591-4d5f-9848-52a9d8b38cab]
[7eea2f69-e31e-4a71-a997-3a370bd7f468]
```

```

Logical_Binding_Stats table
_uuid bytes_from_local bytes_to_local packets_from_local packets_to_local
-----
Logical_Router table
_uuid description name static_routes switch_binding
-----
Logical_Switch table
_uuid description name tunnel_key
-----
-----
be9b786a-d5ab-4f78-8cf7-e15e2c336994 "" "4677abd6-84c1-4aa1-b27b-c06c15eb4b58" 10010
2e70752c-a135-4a94-ab1b-057d51eded0d "" _nvp_internal []
Manager table
_uuid inactivity_probe is_connected max_backoff other_config status target
-----
-----
5cc533b6-3591-4d5f-9848-52a9d8b38cab 30000 false 3000 {} {} "ssl:10.10.50.212:6632"
Mcast_Macs_Local table
MAC _uuid ipaddr locator_set logical_switch
-----
Mcast_Macs_Remote table
MAC _uuid ipaddr locator_set logical_switch
-----
Physical_Locator table
_uuid dst_ip encapsulation_type
-----
e4e82b72-bfb8-4534-8e67-990d92b2e104 "192.168.50.101" "vxlan_over_ipv4"
d69f550e-c90d-4327-a294-465d734a595c "192.168.50.157" "vxlan_over_ipv4"
acb5aea7-db90-436e-a180-93561249c74c "192.168.50.243" "vxlan_over_ipv4"
Physical_Locator_Set table
_uuid locators
-----
Physical_Port table
_uuid description name port_fault_status vlan_bindings vlan_stats
-----
-----
c012b9a0-d840-4de2-9599-499ac27929a0 "" "te-1/1/31" []
{100=be9b786a-d5ab-4f78-8cf7-e15e2c336994} {}
Physical_Switch table
_uuid description management_ips name ports switch_fault_status tunnel_ips tunnels
-----
-----
7eea2f69-e31e-4a71-a997-3a370bd7f468 "" [] "br0" [95d3ce46-6ddd-4897-acae-9417938fc463,
c012b9a0-d840-4de2-9599-499ac27929a0] [] ["192.168.50.101"] []
Tunnel table
_uuid bfd_config_local bfd_config_remote bfd_params bfd_status local remote
-----
-----
Ucast_Macs_Local table
MAC _uuid ipaddr locator logical_switch
-----
-----
"70:72:cf:9d:6f:fb" 50f22ba9-3055-4f68-803b-43bdcc722ae6 ""
d69f550e-c90d-4327-a294-465d734a595c be9b786a-d5ab-4f78-8cf7-e15e2c336994
Ucast_Macs_Remote table
MAC _uuid ipaddr locator logical_switch
-----
-----
"00:50:56:ae:46:d3" 4e8261ea-5b2a-49a3-af28-f5e302fde888 ""
e4e82b72-bfb8-4534-8e67-990d92b2e104 be9b786a-d5ab-4f78-8cf7-e15e2c336994
"00:50:56:ae:5a:30" 6178ea7c-3f90-47fb-bf3f-a7a2f32b0696 ""
e4e82b72-bfb8-4534-8e67-990d92b2e104 be9b786a-d5ab-4f78-8cf7-e15e2c336994

```

OVSDB VTEP with vtep-ctl Configuration Examples

Requirements

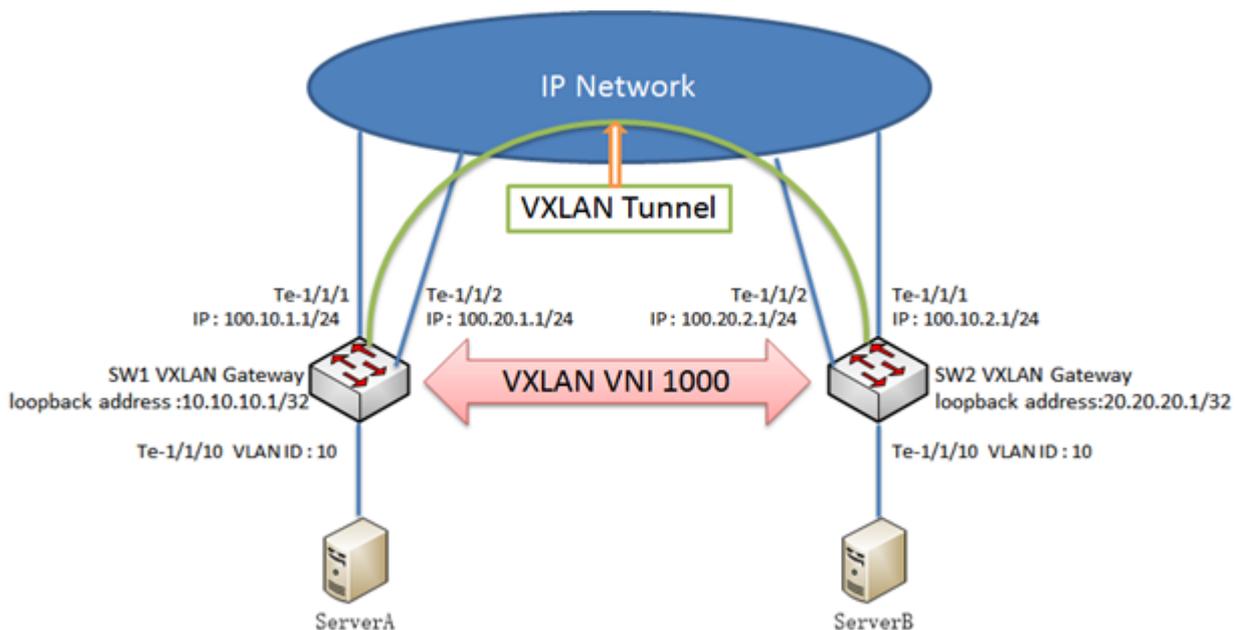
This example uses the following hardware and software components:

- Trident-II Series switch
- PICOS OS Release 2.6 version and higher
- Openvswitch vtep-ctl

Overview

In this example, VTEP OVSDB is configured to run on a VXLAN domain. VTEP interface sources are configured to the loopback address. Interfaces are configured for VLAN tagging and encapsulation. Static route protocols are configured to facilitate unicast routing.

Topology



The VXLAN that comprise the networks include:

- SW1 VNI 1000: VTEP 10.10.10.1: VLAN 10
- SW2 VNI 1000: VTEP 20.20.20.1: VLAN 10
- SW1 IP Adress: 100.10.1.1 and IP Address: 100.20.1.1
- SW2 IP Adress: 100.10.2.1 and IP Address: 100.20.2.1

Configuring OVSDB on Trident-II Series switches

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and copy and paste the commands into the CLI.

SW1 Configure:

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet te-1/1/10 family ethernet-switching vlan members 10
set vlan-interface loopback address 10.10.10.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.1.1 prefix-length 24
set protocols static route 20.20.20.1/32 next-hop 100.10.1.2
set protocols ovssdb controller ovssdb protocol tcp
set protocols ovssdb controller ovssdb address 10.10.50.157
set protocols ovssdb controller ovssdb port 6632
set protocols ovssdb interface te-1/1/10
set vxlan ovssdb-managed true
set vxlans source-interface loopback address 10.10.10.1
```

SW2 Configure:

```
set vlans vlan-id 10
set vlans vlan-id 1000 l3-interface 1000
set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/10 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet te-1/1/10 family ethernet-switching vlan members 10
set vlan-interface loopback address 20.20.20.1 prefix-length 32
set vlan-interface interface 1000 vif 1000 address 100.10.2.1 prefix-length 24
set protocols static route 10.10.10.1/32 next-hop 100.10.2.2
set protocols ovssdb controller ovssdb protocol tcp
set protocols ovssdb controller ovssdb address 10.10.50.159
set protocols ovssdb controller ovssdb port 6632
set protocols ovssdb interface te-1/1/10
set vxlan ovssdb-managed true
set vxlans source-interface loopback address 20.20.20.1
```

Configuring OVSDB Step-by-Step Procedure

The following example shows how to set up a basic VTEP OVSDB configuration with VXLAN domain. To configure VXLAN on an Trident-II Series switches, follow these steps:

Configure OVSDB Step-by-Step for SW1

1. Configure the VLAN ID to 10 for vxlan domain.

set vlans vlan-id 10

2. Configure the VLAN ID to 1000 and 2000 for ip routing.

set vlans vlan-id 1000 I3-interface 1000

3. Configure the te-1/1/1 interface VLAN ID to 1000.

set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000

4. Configure VLAN trunk for te-1/1/10.

set interface gigabit-ethernet te-1/1/10 family ethernet-switching port-mode "trunk"

5. Configure the te-1/1/10 interface VLAN ID to 10.

set interface gigabit-ethernet te-1/1/10 family ethernet-switching vlan members 10

6. Configure ip address for the loopback interface.

set vlan-interface loopback address 10.10.10.1 prefix-length 32

7. Configure ip address for the vlan-interface 1000.

set vlan-interface interface 1000 vif 1000 address 100.10.1.1 prefix-length 24

8. Configure static route for the VXLAN ECMP.

set protocols static route 20.20.20.1/32 next-hop 100.10.1.2

9. Configure VTEP interface sources ip address.

set vxlans source-interface loopback address 10.10.10.1

10. Enable VXLAN managed by ovsdb

set vxlan ovsdb-managed true

11. Configure ovsdb controller protocol

set protocols ovsdb controller ovsdb protocol tcp

12. Configure ovsdb controller ip address

set protocols ovsdb controller ovsdb address 10.10.50.212

13. Configure ovsdb controller port

set protocols ovsdb controller ovsdb port 6632

14. Configure ovsdb controller interface

set protocols ovsdb interface te-1/1/10

Configure OVSDB Step-by-Step for SW2

1. Configure the VLAN ID to 10 for vxlan domain.

set vlans vlan-id 10

2. Configure the VLAN ID to 1000 and 2000 for ip routing.

set vlans vlan-id 1000 I3-interface 1000

3. Configure the te-1/1/1 interface VLAN ID to 1000.

set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 1000

4. Configure VLAN trunk for te-1/1/10.

set interface gigabit-ethernet te-1/1/10 family ethernet-switching port-mode "trunk"

5. Configure the te-1/1/10 interface VLAN ID to 10.

set interface gigabit-ethernet te-1/1/10 family ethernet-switching vlan members 10

6. Configure ip address for the loopback interface.

set vlan-interface loopback address 20.20.20.1 prefix-length 32

7. Configure ip address for the vlan-interface 1000.

```
set vlan-interface interface 1000 vif 1000 address 100.10.2.1 prefix-length 24
```

8. Configure static route for the VXLAN ECMP.

```
set protocols static route 10.10.10.1/32 next-hop 100.10.2.2
```

9. Configure VTEP interface sources ip address.

```
set vxlans source-interface loopback address 20.20.20.1
```

10. Enable VXLAN managed by ovsdb

```
set vxlan ovsdb-managed true
```

11. Configure ovsdb controller protocol

```
set protocols ovsdb controller ovsdb protocol tcp
```

12. Configure ovsdb controller ip address

```
set protocols ovsdb controller ovsdb address 10.10.50.212
```

13. Configure ovsdb controller port

```
set protocols ovsdb controller ovsdb port 6632
```

14. Configure ovsdb controller interface

```
set protocols ovsdb interface te-1/1/10
```

Manual configuration of the Switch as a VTEP Gateway by vtep-ctl

Configuring the ovsdb by vtep-ctl for SW1

1. Create a logical switch

```
vtep-ctl --db=tcp:10.10.51.157:6632 add-ls ls0
```

2. Bind the logical switch to a port

```
vtep-ctl --db=tcp:10.10.51.157:6632 bind-ls br0 te-1/1/10 10 ls0
```

```
vtep-ctl --db=tcp:10.10.51.157:6632 set Logical_Switch ls0 tunnel_key=10010
```

3. Configure Direct unicast destinations out a different tunnel

```
vtep-ctl --db=tcp:10.10.51.157:6632 add-ucast-remote ls0 00:e0:0a:0b:78:9a 20.20.20.1
```

Configuring the ovsdb by vtep-ctl for SW2

1. Create a logical switch

```
vtep-ctl --db=tcp:10.10.51.159:6632 add-ls ls0
```

2. Bind the logical switch to a port

```
vtep-ctl --db=tcp:10.10.51.159:6632 bind-ls br0 te-1/1/10 10 ls0
```

```
vtep-ctl --db=tcp:10.10.51.159:6632 set Logical_Switch ls0 tunnel_key=10010
```

3. Configure Direct unicast destinations out a different tunnel

```
vtep-ctl --db=tcp:10.10.51.159:6632 add-ucast-remote ls0 00:e0:0a:0b:78:9b 10.10.10.1
```

View the VXLAN table of SW1:

```
admin@XorPlus# run show vxlan
```

Egress map:

```
egress_id 100007 MAC 0:0:1:8a:1a:90, port_id 1/1/33, vif_index 6 unicast
```

L3 tunnel mac map:

```
vlan id 1000, ref_count 1
```

Port vlan map mode map & Termination admin state map:

```
port id 1/1/10, ref_count 1
```

Tunnel Map:

```
tunnel id 0X4C000200, dst_vtep 20.20.20.1, nexthops (172.168.1.2 ), ecmp_id 100007, ref_count 1
```

```
tunnel id 0X4C000001, dst_vtep 224.0.0.1, nexthops (), ecmp_id 0, ref_count 1
```

Access ports:

```
id 0X80000002, vpn_id 0X7000, port_id 1/1/10, vlan_id 10, egress_id 100006
```

Network ports:

```
id 0X80000003, vpn_id 0X7000, port_id 1/1/33, egress_id 100007, tunnel_id 0X4C000200, unicast
```

```
id 0X80000004, vpn_id 0X7000, port_id 1/1/33, egress_id 100008, tunnel_id 0X4C000200, multicast
```

```
id 0X80000001, vpn_id 0XFFFFFFFF, port_id 1/1/0, egress_id 100004, tunnel_id 0X4C000001,
multicast
```

BFD sessions:

```
admin@XorPlus#
```

The show vxlan to displays information about VXLAN endpoint configuration. Display VXLAN to the remote vxlan tunnel vtep has two nexthops(100.10.1.2).

View the VXLAN mac table of SW1:

```
admin@XorPlus# run show vxlan ?
```

Possible completions:

```
<[Enter]> Execute this command
```

```
address-table Show vxlan MAC address table
```

```
| Pipe through a command
```

```
admin@XorPlus# run show vxlan address-table
```

```
VNID MAC address Type Interface VTEP
```

```
-----
10010 00:e0:0a:0b:70:00 Static 20.20.20.1
```

L2/L3 Troubleshooting Guide

This guide describes how to identify and resolve common problems related to the Pica8 PicOS software used on supported switches.

- Monitoring and Debugging L2/L3 protocols
- Routing and Forwarding Table
- Using Pipe (|) Filter Functions
- Using the show tech-support Command

Monitoring and Debugging L2/L3 protocols

Find and Configure the Log File

By default, the syslog local-file is ram. The log file name is "message" which is in the directory "/tmp/log"

```
admin@XorPlus$cd /tmp/log
admin@XorPlus$ls
lastlog lighttpd messages wtmp
```

You can use "tail -f /tmp/log/message" to tail the log message.

You can change the syslog local-file to disk. The log file name is "message" which is in the directory of "/var/log"

```
admin@XorPlus# set system syslog local-file disk
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus$cd /var/log/
admin@XorPlus$ls
apt dmesg fsck last_death lastlog messages news ntpstats wtmp
admin@XorPlus$
```

You can use "tail -f /var/log/message" to tail the log message.

Enable Important Debugs

Enable debug interface

```
admin@XorPlus# set interface traceoptions flag ?
Possible completions:
<[Enter]> Execute this command
all Configure all tracing
config Configure configuration tracing
ethernet-switching-options Configure ethernet-switching-options tracing
mlog-trace Configure mlog event tracing
neighbor-event Configure neighbor event tracing
```

```

packets Configure received or sent packets event tracing
port-security Configure port security tracing
raw-packet Configure receive raw packet tracing
route-event Configure route event tracing
static-ethernet-switching Configure static-ethernet-switching tracing
admin@XorPlus# set interface traceoptions flag config ?
Possible completions:
<[Enter]> Execute this command
disable Disable tracing
admin@XorPlus# set interface traceoptions flag config disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#
admin@XorPlus# set interface traceoptions line-card ?
Possible completions:
<[Enter]> Execute this command
statistics Configure line card statistic module trace
trace-level Configure line card trace level
trace-type Configure line card trace type
admin@XorPlus# set interface traceoptions line-card trace-level all disable false
admin@XorPlus# commit
Commit OK.

```

Enable debug of BGP:

```

admin@XorPlus# set protocols bgp traceoptions flag ?
Possible completions:
<[Enter]> Execute this command
FSM Configure FSM state tracing
all Configure all tracing operations
configuration Configure configuration tracing
damp Configure damp tracing
event Configure FSM Event tracing
keepalive Configure keepalive packet tracing
nexthop-resolution Configure nexthop resolution with RIB tracing
notification Configure notification packet tracing
notify-snmp Configure notify snmp tracing
open Configure open packet tracing
packet Configure all packet tracing operations
policy Configure policy configuration tracing
raw-packet Dump raw packets BGP read from or send to sockets
route-refresh Configure route refresh packet tracing
timer Configure timers tracing
update Configure update packet tracing
admin@XorPlus# set protocols bgp traceoptions flag all disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

Enable debug of ospf4:

```

admin@XorPlus# set protocols ospf4 traceoptions flag ?
Possible completions:
<[Enter]> Execute this command
adjacency-event Configure tracing of adjacency events
all Configure all tracing operations
config Configure tracing of configuration
database-description Configure tracing of database description packets
event Configure tracing of OSPF state machine events
flooding Configure tracing of lsa flooding
hello Configure tracing of hello packets
lsa-ack Configure tracing of LSA acknowledgment packets
lsa-generation Configure tracing of lsa-generation events

```

```

lsa-request Configure tracing of LSA request packets
lsa-update Configure tracing of LSA update packets
packets Configure tracing of all OSPF packets
retransmission Configure tracing of lsa-packets retransmission
route Configure tracing of routing information
spt Configure tracing of spt calculations
timer Configure tracing of routing protocol timer processing
admin@XorPlus# set protocols ospf4 traceoptions flag all disable false
admin@XorPlus# commit

```

Enable debug of stp:

```

admin@XorPlus# set protocols spanning-tree traceoptions interface ge-1/1/1 ?
Possible completions:
<[Enter]> Execute this command
all Configure all tracing operations
bridge-detection-machine Configure bridge detection state machine tracing
configuration Configure configuration tracing
events Configure events tracing
message-in Configure receive message tracing
message-out Configure send message tracing
port-information-machine Configure port information state machine tracing
port-migration-machine Configure port migration state machine tracing
port-receive-machine Configure port receive state machine tracing
port-role-selection-machine Configure port role selection state machine tracing
port-role-transition-machine Configure port role transition state machine tracing
port-state-transition-machine Configure port state transition state machine tracing
port-transmit-machine Configure port transmit state machine tracing
state-machine-variables Configure state machine variables tracing
timers Configure timers tracing
topology-change-machine Configure topology change state machine tracing
admin@XorPlus# set protocols spanning-tree traceoptions interface ge-1/1/1 all disable
false
admin@XorPlus# set protocols spanning-tree traceoptions interface ge-1/1/1 all disable
false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

Enable debug of igmp:

```

admin@XorPlus# set protocols igmp traceoptions flag ?
Possible completions:
<[Enter]> Execute this command
all Configure all tracing operations
event Configure event tracing
leave Configure leave tracing
query Configure query tracing
report Configure report tracing
admin@XorPlus# set protocols igmp traceoptions flag all disable false
admin@XorPlus# commit
Commit OK.
Save done.
admin@XorPlus#

```

Find the Core Dump File

When the device crashes, it will create a core file which can be found in a directory called *pica/core*.

```
admin@XorPlus$pwd
/pica/core
```

Find the last_death file for Troubleshooting

You can view the last_death file after the device crashes. It will record the last log message.

```
admin@XorPlus$pwd
/var/log
admin@XorPlus$ls
apt dmesg fsck last_death lastlog messages news ntpstats wtmp
```

Routing and Forwarding Table

Check the Software and Hardware Route Tables

To display the hardware *host* route table, use the **show route forward-host ipv4 all** command in L2/L3 operation mode.

```
admin@Switch> show route forward-host ipv4 all
Address HWaddress Port
-----
10.10.3.2 48:6E:73:02:03:DA ge-1/1/48
Total host count:1
```

To display the hardware route table, use the **show route forward-route ipv4 all** command in L2/L3 operation mode.

```
admin@Switch> show route forward-route ipv4 all
Destination          NextHopMac          Port
-----
10.10.3.0/24         48:6E:73:02:04:64  connected
101.101.101.0/24     48:6E:73:02:03:DA  ge-1/1/48
102.102.102.0/24     48:6E:73:02:03:DA  ge-1/1/48
Total route count:4
```

To display the software route table, use the **show route table ipv4 unicast static** command in L2/L3 operation mode.

```
admin@Switch> show route table ipv4 unicast static
```

If PicOS is running in OVS mode, check the software and hardware flow tables.

Pipe (|) Filter Functions

This topic describes the pipe (|) filter functions supported in the Pica8 PicOS L2/L3 CLI (command-line interface). The PicOS L2/L3 mode has a growing number of CLI commands that users can use to troubleshoot common problems. These commands usually generate a lot of output. The use of pipe (|) filter functions increases readability of command output, making troubleshooting more effective.

The following filter functions are available with the PicOS L2/L3:

Function	Description
compare	Compare configuration changes with a prior version
count	Count occurrences
display	Display additional configuration information
except	Show only the lines of output that do not contain a pattern
find	Show output starting from the first occurrence of a pattern
match	Show only the lines of output that contain a pattern
no-more	Disable pagination of command output

Comparing Configurations

The **compare** filter compares the current committed configuration with a previously committed configuration.

```
admin@XorPlus# show | compare rollback nn
```

nn is the index into the list of previously committed configurations, also known as the rollback number. The range of values for **nn** is 01-48.

For example:

```
admin@XorPlus# show | compare rollback 03
```

Counting the Number of Output Lines

To count the number of lines in the output of a command, enter **count** after the pipe symbol (|).

The following example uses **count** with the **show** command in configuration mode to display the number of non-default configuration lines:

```
admin@XorPlus# show | count
Count: 11 lines
```

Displaying Output that Matches a Pattern

To display only the lines of output that match a pattern, enter **match** after the pipe symbol (|).

The following example displays the status of only TbE (terabit Ethernet) interfaces:

```
admin@XorPlus> show interface brief | match te-
te-1/1/1    Enabled    Down    Disabled    Full    Auto
```

```

te-1/1/2    Enabled    Down    Disabled    Full    Auto
te-1/1/3    Enabled    Down    Disabled    Full    Auto
te-1/1/4    Enabled    Down    Disabled    Full    Auto
te-1/1/5    Enabled    Down    Disabled    Full    Auto
te-1/1/6    Enabled    Down    Disabled    Full    Auto
te-1/1/7    Enabled    Down    Disabled    Full    Auto
te-1/1/8    Enabled    Down    Disabled    Full    Auto
te-1/1/9    Enabled    Down    Disabled    Full    Auto
te-1/1/10   Enabled    Down    Disabled    Full    Auto
<Some output omitted>

```

Omitting Output that Matches a Pattern

To omit lines from the output of a command that make up a pattern, enter **except** after the pipe symbol (`|`).

The following example uses **except** with the **show interface brief** command in the operation mode to list the interfaces that are *not* down:

```
admin@XorPlus> show interface brief | except Down
```

Preventing Output from Being Paginated

By default, if the output of a command is longer than the length of terminal screen, user will see the **--More--** message to display the remaining output. Press the space bar to display the remaining output.

User can disable pagination by entering **no-more** after the pipe symbol (`|`).

The following example displays the output of **show** command, executed in PicOS L2/L3 configuration mode, all at once:

```
admin@XorPlus# show | no-more
```

This feature is useful, for example, when user wants to copy the entire output of a command and paste it into an e-mail to be sent to technical support.

Show Tech-Support Command

When contacting Pica8 for technical support, issue the command **show tech-support** because it captures the complete status of a PicOS switch. It is recommended to send the output of **show tech-support** command along with the system log. The following samples describe how to obtain the output.

Log in to the switch and enter the **cli** command at the Linux shell to reach the PicOS L2/L3 operation mode.

```

admin@Leaf-1$cli
Synchronizing configuration...OK.
Pica8 PicOS Version 2.6
Welcome to PicOS L2/L3 on Leaf-1
admin@Leaf-1>

```

Enter the **show tech-support** command.

```
admin@Leaf-1> show tech_support
Start.....

Item 1: Display system version finished!
Item 2: Display system interface finished!
Item 3: Display system configuration finished!
Item 4: Display system config files finished!
Item 5: Display system process finished!
Item 6: Display system fdb table finished!
Item 7: Display system fdb entries finished!
Item 8: Display system ospf neighbors finished!
Item 9: Display system ospf interfaces finished!
Item 10: Display system route table finished!
Item 11: Get error event from log!
Item 12: Display system hard-route table finished!
Item 13: Display system hard-route for host finished!
Item 14: Display system spanning tree interfaces finished!
Item 15: Display system spanning tree bridge finished!
Item 16: Display system vlans table finished!
Item 17: Display system vlan-interfaces finished!
Item 18: Display system core-dump finished!
Item 19: Display system uptime finished!
Item 20: Display system arp table finished!

The information has been stored in /tmp/Leaf-1-201507050614-techSupport.log, please forward
to support@pica8.com
```

The last line of the output of **show tech-support** command provides the name and location of the file to which the output was saved. In the above example, the name of the file is *Leaf-1-201507050614-techSupport.log* that has been saved to the */tmp* directory.

You can transfer the file, generated by **show tech-support** command, from the switch to your computer over SCP (Secure Copy Protocol). There is a nice free Windows utility called WinSCP, available for download at <https://winscp.net/eng/download.php>, which you can use to copy the file from the switch to your computer over SCP.

Configuration Appendix

- Other Command List

Other Command List

```
set interface traceoptions flag config disable true
set interface traceoptions flag ethernet-switching-options disable true
set interface traceoptions flag mlag-trace disable true
set interface traceoptions flag neighbor-event disable true
set interface traceoptions flag packets disable true
set interface traceoptions flag route-event disable true
set interface traceoptions flag static-ethernet-switching disable true
set interface traceoptions line-card statistic disable true
set interface traceoptions line-card trace-level all disable true
set interface traceoptions line-card trace-level api debug disable true
set interface traceoptions line-card trace-level api error disable true
set interface traceoptions line-card trace-level api information disable true
set interface traceoptions line-card trace-level api warning disable true
set interface traceoptions line-card trace-level sdk debug disable true
set interface traceoptions line-card trace-level sdk error disable true
set interface traceoptions line-card trace-level sdk information disable true
set interface traceoptions line-card trace-level sdk warning disable true
set interface traceoptions line-card trace-level xrl debug disable true
set interface traceoptions line-card trace-level xrl error disable true
set interface traceoptions line-card trace-level xrl information disable true
set interface traceoptions line-card trace-level xrl warning disable true
set interface traceoptions line-card trace-type all disable true
set interface traceoptions line-card trace-type configuration disable true
set interface traceoptions line-card trace-type link-change disable true
set interface traceoptions line-card trace-type mac-update disable true
set interface traceoptions line-card trace-type packet disable true
set interface traceoptions line-card trace-type packet-receive disable true
set interface traceoptions line-card trace-type packet-transmit disable true
set interface traceoptions line-card trace-type statistics disable true
```