



CHAPTER 3

NVT Summary, Issues, and Recommendations

This chapter contains the following sections:

- [Summary, page 3-1](#)
- [Issues and Recommendations, page 3-1](#)
- [Conclusion, page 3-6](#)

Summary

The test results verify that the services and features configured and tested (as described in this document) at each network layer of a running Cisco Catalyst platform can be deployed using similar feature sets on the Cisco Nexus 7000. In general, the results show that after any of the disruptions and subsequent recoveries, the network reconverges to the expected state within the expected time frame.

Issues and Recommendations

Specific issues and recommendations derived from the testing experience are discussed in this section. This section contains the following topics:

- [Optimization of CoPP, page 3-2](#)
- [BFD or Aggressive Timers, page 3-3](#)
- [ISSU, page 3-3](#)
- [Routing Protocols and MTU, page 3-4](#)
- [Reserved VLANs, page 3-4](#)
- [Link Aggregation Protocols, page 3-5](#)
- [Multichassis EtherChannel Interoperability, page 3-5](#)
- [IS-IS Default Metric Style, page 3-5](#)
- [OSPF Interface Metric, page 3-5](#)
- [BGP Per-Interface Fast External Failover, page 3-5](#)
- [Multicast Multipath, page 3-5](#)
- [Multicast SPT Threshold, page 3-6](#)
- [Dynamic Trunking Protocol, page 3-6](#)

Optimization of CoPP

CoPP optimization is recommended on the Nexus 7000 to achieve expected multicast routing scale and performance on the first hop source router for ASM. By default, any directly-connected multicast source will have its data traffic rate-limited by the Layer 3 multicast directly-connected rate limiter and further policed by the CoPP class default. This double layer of limiting may affect PIM source registration performance, especially when a large number of sources come online at the same time.

For NVT, the Layer 3 multicast directly-connected rate limiter is disabled. A new CoPP class is created to police multicast source data traffic so that multicast source registration performance with the NVT test profile is comparable between Nexus 7000 and Catalyst 6500 Supervisor Engine 2T. The ACL used in this CoPP class ensures that control plane protocol packets are not policed.

In order to update the CoPP configuration on the Nexus 7000, enter the following command to create a copy of the default configuration:

```
copp copy profile lenient prefix test
```

Enter the following commands to apply the copy to the control plane interface:

```
control-plane
  service-policy input test-copp-policy-lenient

hardware rate-limiter layer-3 multicast directly-connected disable
ip access-list multicast-source-data
  10 deny ip any 224.0.0.0/24
  20 deny ip any 224.0.1.0/24
  30 permit ip any 224.0.0.0/4

class-map type control-plane match-any multicast-source-data
  match access-group name multicast-source-data

policy-map type control-plane test-copp-policy-lenient
  class test-copp-class-critical
    set cos 7
    police cir 39600 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-important
    set cos 6
    police cir 1060 kbps bc 1500 ms conform transmit violate drop
  class test-copp-class-management
    set cos 2
    police cir 10000 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-normal
    set cos 1
    police cir 680 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-normal-dhcp
    set cos 1
    police cir 680 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-normal-dhcp-relay-response
    set cos 1
    police cir 900 kbps bc 750 ms conform transmit violate drop
  class test-copp-class-redirect
    set cos 1
    police cir 280 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-exception
    set cos 1
    police cir 360 kbps bc 375 ms conform transmit violate drop
  class test-copp-class-monitoring
    set cos 1
    police cir 130 kbps bc 1500 ms conform transmit violate drop
  class test-copp-class-l2-unpoliced
    police cir 8 gbps bc 5 mbytes conform transmit violate transmit
```

```

class test-copp-class-undesirable
  set cos 0
  police cir 32 kbps bc 375 ms conform drop violate drop
class test-copp-class-l2-default
  police cir 100 kbps bc 375 ms conform transmit violate drop
class multicast-source-data
  police cir 1000 kbps bc 250 ms conform transmit violate drop
class class-default
  set cos 0
  police cir 100 kbps bc 250 ms conform transmit violate drop

```

**Note**

On the Catalyst 6500 Supervisor Engine 720, the following rate limiters must be configured. In the absence of rate limiters, control plane protocols like BFD with short keepalive intervals or aggressive timers may flap.

```

mls rate-limit multicast ipv4 fib-miss 1000 100
mls rate-limit multicast ipv4 non-rpf 1000 100
mls rate-limit multicast ipv4 connected 1000 100
mls rate-limit multicast ipv4 partial 1000 100
mls rate-limit multicast ipv6 connected 1000 100
mls rate-limit multicast ipv6 mld 10 100

```

On the Catalyst 6500 Supervisor Engine 2T, there is no need for additional configuration to protect the control plane.

BFD or Aggressive Timers

BFD is recommended due to lower control plane overhead in order to achieve fast network failure detection and reconvergence. With BFD, any number of supported clients can piggy-back on top of one BFD session per connection for fast failure detection. The Nexus 7000 further enhances this capability with distributed BFD where BFD runs per line card instead of on the supervisor. The list of clients supporting BFD on the Nexus 7000 is extensive; however, the Catalyst 6500 and 4500 do not support BFD on all types of interfaces, and the list of BFD clients supported is less extensive. Some common examples of unsupported interfaces are port-channels, SVIs, and sub-interfaces; some examples of unsupported clients are the FHRP protocols and PIM. For connections between the Nexus 7000 and these platforms requiring fast peer failure detection for unsupported interfaces and unsupported BFD clients, aggressive timers are used.

The BFD retransmit interval is configured to be 1 second with 3x holddown multiplier. These parameters are chosen to match the protocols running aggressive timers with 1 second hello intervals.

ISSU

ISSU and rollback was performed between 5.2.5 and the following images:

- 5.2.4
- 5.2.3a
- 5.1.6
- 6.1.1

Upgrades and the rollbacks were tested with the ‘parallel’ option (where applicable) to minimize network maintenance window.

For networks running EIGRP, ISSU may cause routing peers to flap, especially for very high-scaled networks.

For switches running VTP in server mode, ISSU is recommended. Otherwise, an image version change via reload may cause VTP VLAN configuration to be lost.

Routing Protocols and MTU

If jumbo MTU is configured on the Nexus 7000, unicast routing protocols will leverage jumbo MTU for control plane update packets. PIM does not use jumbo MTU.

IOS on the Catalyst 6500 does not use jumbo MTU to send control plane protocol updates.

This is for information only. No additional configuration is required for the systems to fully interoperate.

Reserved VLANs

On the Catalyst 6500 Supervisor Engine 2T and the Nexus 7000, VLANs are not reserved for Layer 3 LAN ports and subinterfaces; however, some software features use internal VLANs in the extended range.

On the Catalyst 6500 Supervisor Engine 720 and Catalyst 4500/4948, VLANs are reserved for Layer 3 LAN ports and subinterfaces, and some software features use internal VLANs in the extended range.

You cannot use any VLAN that has been allocated for internal use.

For additional details, see the following references:

NX-OS 5.x releases

http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/nx-os/layer2/configuration/guide/Cisco_Nexus_7000_Series_NX-OS_Layer_2_Switching_Configuration_Guide_Release_5.x_chapter4.html#task_67E5266F50104AF38E5149C1CC56B1A7

Catalyst 6500 SX releases

<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/vlans.html#wpmkr1037585>

Catalyst 6500 Sup2T 15.0SY releases

<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.0SY/configuration/guide/vlans.html#wpmkr1037585>

Catalyst 4500 XE 3.3.0

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/15.1/XE_330SG/configuration/guide/vlans.html

Catalyst 4900 12.2(46)SG

<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/46sg/configuration/guide/vlans.html>

Link Aggregation Protocols

IOS—PAgP (auto/desirable) is supported.

NX-OS— PAgP (auto/desirable) is not supported.

**Note**

NX-OS does not support PAgP. When a Catalyst 6500 VSS is connected to a Nexus 7000, enhanced PAgP for VSS Dual-Active Detection cannot be used on that link.

Multichassis EtherChannel Interoperability

To reduce data loss following a stateful switchover (SSO) on a Catalyst 6500 VSS, port load-share deferral is recommended on a port channel of a switch that is connected by a multichassis EtherChannel (MEC) to a VSS. Port load-share deferral is not supported on the Catalyst 4500/4948 and the Nexus 7000.

IS-IS Default Metric Style

On IOS systems, IS-IS must be configured to use transit or wide metric style to interoperate with the Nexus 7000.

IOS—Defaults to old style (narrow).

NX-OS—Defaults to new style (wide).

OSPF Interface Metric

The reference bandwidth for calculating OSPF metric is 100Mbps for IOS and 40Gbps for NX-OS. The entire OSPF network should be configured to use the same reference bandwidth. For NVT, 100Gbps was used.

BGP Per-Interface Fast External Failover

On NX-OS, per-interface fast failover applies only to eBGP peers. On IOS, it applies to all BGP peer types.

Multicast Multipath

IOS—Multicast multipath is disabled by default. When multipath is enabled, the default load sharing selection algorithm is source-based. The algorithm on IOS can be configured to match the behavior on NX-OS with the following command:

```
ip multicast multipath s-g-hash basic
```

NX-OS—Multicast multipath is enabled by default and the load sharing selection algorithm is based on the source and group addresses.

Multicast SPT Threshold

IOS—Multicast group filtering for spt-threshold is configured using an IP access list.

NX-OS—Multicast group filtering for spt-threshold is configured using a route-map. Within the route-map, the group filter can be specified using prefix address and mask; the group-range command, though available, is not supported.

Dynamic Trunking Protocol

Dynamic Trunking Protocol (DTP) is not supported in NX-OS. Configure the trunk port for unconditional trunking on the Catalyst 6500 when interoperating with the Nexus 7000.

Conclusion

By following the recommendations and guidelines suggested in this document, customers who deploy Cisco Catalyst 6500 and Nexus 7000 Series Switches can expect that the hardware and software features used in this network will behave and perform similarly between the platforms. Since NVT has been established as an additional quality assurance stage in order to leverage customer feedback and requirements into the product development cycle, future phases of NVT will continue to validate and publish additional guidelines for deploying Nexus 7000 and NX-OS solutions for datacenter networks. This document is intended to supplement the Cisco Nexus 7000 Series Switches product documentation that is available on cisco.com and should not be used as a replacement for that documentation.