networktest

Juniper / Cisco Interoperability Cookbook

August 2014

TABLE OF CONTENTS

ILLUSTRATIONS

Introduction

Objectives

This configuration guide aims to help networking professionals successfully interconnect Juniper Networks and Cisco Systems switches using a variety of popular Layer 2 and Layer 3 protocols. By following the step-by-step procedures described in this document, it should be possible to verify interoperability and to pass traffic between the two vendors' switches.

Intended audience

This configuration guide is intended for any network architect, administrator, or engineer who needs to interconnect Juniper and Cisco Ethernet switches.

This document assumes familiarity with basic Ethernet and TCP/IP networking concepts, as well as at least limited experience with the Juniper and Cisco command-line interfaces (CLIs). No previous experience is assumed for the protocols discussed in this document.

For beginning readers unfamiliar with Juniper or Cisco CLI syntax, both companies' web sites offer free access to extensive software documentation. In addition, several excellent books on Juniper Junos Software and Cisco IOS configuration are available.

For Juniper Junos operating system configuration, these titles include *Junos Enterprise Switching* by Harry Reynolds and Doug Marschk; *Day One: Exploring the Junos CLI* by Cathy Gadecki and Michael Scruggs, available in free PDF format or in book format; and the widely used *Junos Cookbook* by Aviva Garrett.

Popular titles on Cisco IOS configuration include *Cisco LAN Switching Fundamentals* by David Barnes and Basir Sakandar; *Cisco Routers for the Desperate* by Michael W. Lucas; and *Routing TCP/IP, Volume 1* by Jeff Doyle and Jennifer Carroll.

For basic TCP/IP networking concepts, the standard references are *Internetworking with TCP/IP, Volume 1* by Douglas E. Comer and *TCP/IP Illustrated, Volume 1* by Kevin R. Fall and W. Richard Stevens.

For IP multicast topics, *Interdomain Multicast Routing: Practical Juniper Networks and Cisco Systems Solutions* by Brian M. Edwards, Leonard A. Giuliano, and Brian R. Wright offers in-depth explanations of multicast routing protocols and numerous configuration examples using Juniper and Cisco routers.

Devices covered in this document

Using the commands given in this document, Network Test has verified interoperability between the Juniper EX4300, QFX5100, and Juniper EX9200 Ethernet switches and Cisco Catalyst 3850 and Cisco Nexus 7000 series Ethernet switches. The Layer-3 VPN

interoperability section uses a Juniper MX80 router as well as the other devices previously mentioned. The Wi-Fi interoperability section also uses a Cisco 5508 controller and Cisco 3602 and Cisco 3702 access points. Appendix B lists software versions tested. Except where specifically noted, command syntax for the Juniper and Cisco switches does not change across product lines.

Conventions used in this document

The typographical syntax in this document follows that used in the Juniper *Complete Software Guide for Junos Software for EX Switches.*

The following table lists text and syntax conventions.

Convention	Description	Examples
Bold type	Represents text that you type	To enter configuration mode, type the **configure** command: `admin@host> configure`
Fixed-width text like this	Represents output that appears on the terminal screen	`admin@host> show chassis alarms` `No alarms currently active`
Italic text like this	• Introduces important new terms • Identifies book titles • Identifies RFC and Internet-draft titles • Identifies variables (options for which you substitute a value) in commands or configuration statements.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos System Basics Configuration Guide* • RFC 4814, *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking* • `admin@# set system domain-name domain-name`
< > angle brackets	Enclose optional keywords or variables.	`stub <default-metric metric>;`
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	`broadcast | multicast (string1 | string2 | string3)`
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it appears.	`rsvp { # Required for dynamic MPLS only`
[ ] (square braces)	Enclose a variable for which you can substitute one or more values.	`community name members [ community-ids]`

Indention and braces ({ })	Identify a level in the configuration hierarchy.	`[edit]`
		`routing-options {`
		`  static {`
		`  route default {`
		`  nexthop address;`
		`  retain;`
		`  }`
		`  }`
		`}`

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	`nexthop address;`

Interoperability testing

For each interoperability test described here, this document uses a five-section format consisting of objective, technical background, Juniper configuration, Cisco configuration and test validation.

Cisco Discovery Protocol (CDP) passthrough

Objective

To verify the ability of a Juniper switch to forward Cisco Discovery Protocol (CDP) traffic between two Cisco devices.

Background

The proprietary Cisco Discovery Protocol (CDP) allows sharing of information, such as IP address, model number, and power requirements among connected Cisco devices. Cisco devices use CDP messages to transmit information about their capabilities to other Cisco products in the network. Accordingly, an interoperability requirement for any Juniper switch in the path between two Cisco devices is the ability to "pass through" CDP traffic without affecting CDP operation.

No extra configuration of Juniper or Cisco switches is required for CDP passthrough. Because Juniper EX Series and QFX Series switches forward CDP messages in regular Ethernet frames, a standard Ethernet switching configuration will work. Similarly, CDP is enabled by default on most Cisco devices, so no additional configuration is needed.

Topology

In this example, Cisco Catalyst 3850 and Cisco Nexus 7010 switches will use CDP to exchange model numbers and interface information across two Juniper EX9208 switches in a Virtual Chassis configuration. Though not required for this test, the inter-switch links also used link aggregation to bundle one or more physical interfaces into a single logical pipe. There is a separate section in this document describing link aggregation configuration.

The interfaces used are as follows:

- Cisco Catalyst 3850: TenGigabitEthernet1/1/3, TenGigabitEthernet1/1/4, and Port-channel2 (t1/1/3, t1/1/4 and po2)
- Juniper Virtual Chassis with EX9208: xe-5/3/1, xe-12/3/0, and ae1 (to Catalyst 3850); and xe-5/0/5, xe-12/0/5, and ae2 (to Nexus 7010)
- Cisco Nexus 7010: Ethernet3/9, Ethernet3/10, and port-channel1 (e3/9, e3/10, and po1)

All devices are configured as switches and all inter-switch links act as VLAN trunks. However, even without VLAN configuration CDP traffic will be forwarded just as in this example. Figure 1 shows the topology for CDP passthrough.

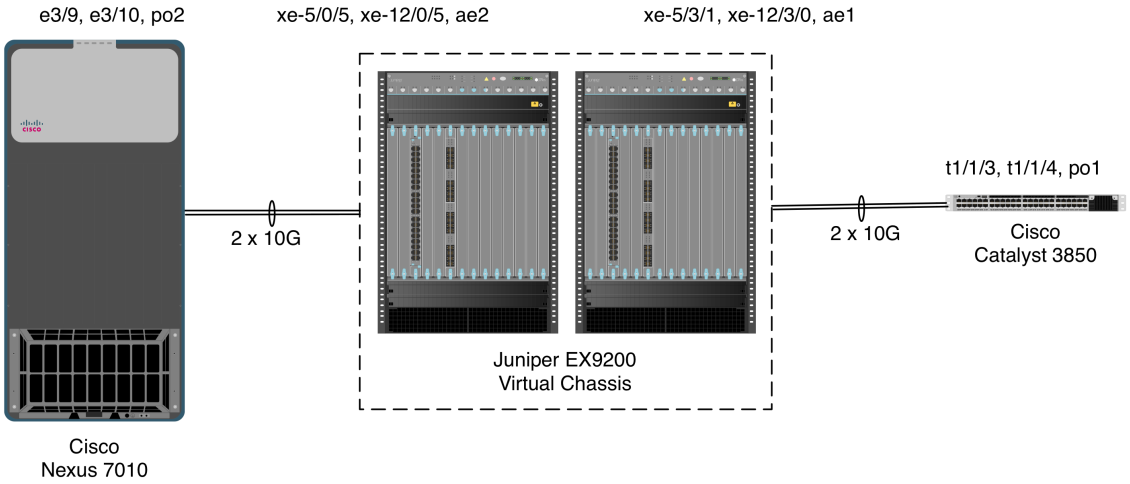

Figure 1: CDP passthrough topology

Juniper configuration

In this example, the inter-switch ports use link aggregation and VLAN trunking. Both steps are optional; if desired, physical switch ports can instead be configured for Ethernet switching.

These steps define two VLANs; again, this is optional:

```
admin@EX9208> configure
admin@EX9208# set vlans v2001 vlan-id 2001
admin@EX9208# set vlans v2002 vlan-id 2002
```

These steps create link aggregation groups **ae1** to the Cisco Catalyst 3850 and **ae2** to the Cisco Nexus 7010:

```
admin@EX9208# set interfaces ae1 aggregated-ether-options lacp active
admin@EX9208# set interfaces ae1 description "linkagg to 3850"
admin@EX9208# set interfaces ae2 aggregated-ether-options lacp active
admin@EX9208# set interfaces ae2 description “linkagg to 7010"
```

These steps configure the newly created **ae** interfaces as Ethernet switching and VLAN trunk ports:

```
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching
interface-mode trunk
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching vlan
members v2001
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching vlan
members v2002
```

```
admin@EX9208# set interfaces ae2 unit 0 family ethernet-switching
interface-mode trunk
admin@EX9208# set interfaces ae2 unit 0 family ethernet-switching vlan
members v2001
admin@EX9208# set interfaces ae2 unit 0 family ethernet-switching vlan
members v2002
```

These steps assign physical ports to membership in the link aggregation groups:

```
admin@EX9208# set interfaces xe-5/3/1 ether-options 802.3ad ae1
admin@EX9208# set interfaces xe-5/3/1 description "ae1 linkagg to 3850"
admin@EX9208# set interfaces xe-12/3/0 ether-options 802.3ad ae1
admin@EX9208# set interfaces xe-12/3/0 description "ae1 linkagg to 3850"
admin@EX9208# set interfaces xe-5/0/5 ether-options 802.3ad ae2
admin@EX9208# set interfaces xe-5/0/5 description "ae2 linkagg to 7010"
admin@EX9208# set interfaces xe-12/0/5 ether-options 802.3ad ae2
admin@EX9208# set interfaces xe-12/0/5 description "ae2 linkagg to 7010"
```

The spanning tree protocol must be either disabled on all switches, or disabled on all switches. This command will enable rapid spanning tree on a Juniper EX Series switch:

```
admin@EX9208# set protocols rstp
admin@EX9208# commit
```

To disable rapid spanning tree on a Juniper EX Series switch:

```
admin@EX9208# set protocols rstp disable
admin@EX9208# commit
```

Cisco commands

Since CDP is enabled by default on Cisco devices, no additional configuration is needed. The steps given here are to define link aggregation and VLAN trunking, but both are optional for purposes of validating CDP passthrough.

On the Catalyst 3850, CDP and rapid spanning tree (called Rapid PVST-Plus in Cisco documentation) will be enabled. All that remains is to (optionally) define a link aggregation group and VLAN trunking.

These commands will create two VLANs:

```
Cat3850# configure terminal
Cat3850(config)# vlan 2001-2002
Cat3850(config-vlan)# exit
```

These steps will create a link aggregation group (called a Port-channel in Cisco parlance) and configure it for VLAN trunking:

```
Cat3850(config)# interface Port-channel1
Cat3850(config-if)# description to EX9200
Cat3850(config-if)# switchport trunk allowed vlan 2001-2002
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# exit
```

These steps assign physical ports to membership in the link aggregation groups:

```
Cat3850(config)# interface TenGigabitEthernet1/1/3
Cat3850(config-if)# description po1 to 9200
Cat3850(config-if)# switchport trunk allowed vlan 2001-2002
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# channel-group 1 mode passive
Cat3850(config)# interface TenGigabitEthernet1/1/4
Cat3850(config-if)# description po1 to 9200
Cat3850(config-if)# switchport trunk allowed vlan 2001-2002
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# channel-group 1 mode passive
Cat3850(config-if)# end
```

Then issue similar commands on the Nexus 7010. First create multiple VLANs:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001-2002
Nexus7010(config-vlan)# exit
```

Then create a link aggregation group and configure it for VLAN trunking:

```
Nexus7010(config)# interface port-channel2
Nexus7010(config-if)# description linkagg to ex9200 ae2
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2002
```

Finally, assign physical ports to the link aggregation group and configure it for VLAN trunking. Note that Cisco Nexus ports are in **shutdown** mode by default, and must be explicitly enabled:

```
Nexus7010(config-if)# interface Ethernet3/9
Nexus7010(config-if)# description linkagg to ex9200 ae2 xe-5/0/5 and
xe-12/0/5
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# channel-group 2 mode active
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Validation

To verify that a Juniper EX Series switch will forward CDP messages between two Cisco devices, use the **show cdp neighbors** command on either Cisco device.

The Catalyst 3850 will recognize the Nexus 7010 via CDP:

```
Cat3850#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Nexus7010(TBM13093202)
                 Ten 1/1/4         138             R S C  N7K-C7010 Eth
3/10
Nexus7010(TBM13093202)
                 Ten 1/1/3         138             R S C  N7K-C7010 Eth
3/10
```

And the Nexus 7010 similarly will recognize the Catalyst 3850:

```
Nexus7010# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform      Port ID
Cat3850.englab.juniper.net
                    Eth3/9         156    S I       WS-C3850-48P  Ten1/1/3
Cat3850.englab.juniper.net
                    Eth3/9         170    S I       WS-C3850-48P  Ten1/1/4
```

Note that in both cases, the Cisco switches correctly identified the hostname ("Device ID"), model number ("Platform") and interface ("Port ID") of the remote Cisco device. All this information is learned via CDP, which is forwarded without any additional configuration needed on Juniper switches running Junos.

Generic Routing Encapsulation (GRE)

Objective

To verify the ability of Juniper and Cisco switches to tunnel traffic over an IP backbone using GRE.

Background

As described in IETF RFC 2784, GRE provides a method of encapsulating traffic into IP packets for transmission across a routed network. On the receiving end, the traffic is decapsulated and forwarded in its original form.

Although this configuration example uses IP-in-IP encapsulation, GRE can carry virtually any protocol, including non-routable traffic such as raw Ethernet frames, across a routed IP network.

Topology

Figure 2 shows the GRE validation test bed. This example uses a GRE tunnel between a Juniper Virtual Chassis (here, comprising two Juniper EX9208 switches) and a Cisco Nexus 7010 switch. The GRE tunnel endpoints are in the 192.18.44.0/24 subnet, and the tunnel uses a 10-gigabit Ethernet link between switches.

Even though a different subnet (192.18.38.0/24) is configured on the physical interfaces of both switches, traffic will traverse the GRE tunnel. In this example, the Juniper and Cisco devices use OSPF to learn about network reachability.

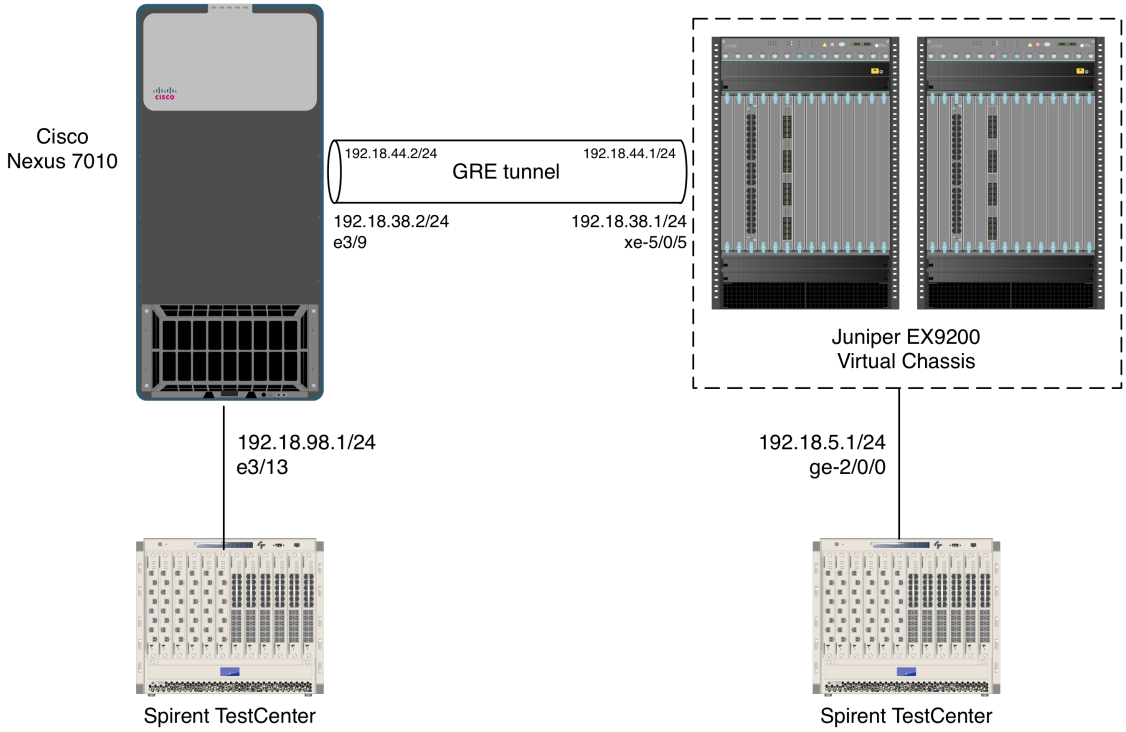

Figure 2: GRE validation topology

Juniper commands

1. Define IP addresses on the inter-switch link and the link to the Spirent TestCenter traffic generator/analyzer. Note that these addresses are associated with the physical interfaces, not the GRE tunnel:

```
admin@ex9208> configure
admin@ex9208# set interfaces xe-5/0/5 description "physical interface to
Nexus 7010 e3/9"
admin@ex9208# set interfaces xe-5/0/5 unit 0 family inet address
192.18.38.1/24
admin@ex9208# set interfaces ge-2/0/0 description "physical interface to
STC"
```

```
admin@ex9208# set interfaces ge-2/0/0 unit 0 family inet address
192.18.5.1/24
```

2. Define a GRE tunnel, in this case called **gr-5/0/0**. In Junos, the required commands include a tunnel source, destination, and endpoint IP address:

```
admin@ex9208# set interfaces gr-5/0/0 description "GRE tunnel endpoint to
Nexus 7010"
admin@ex9208# set interfaces gr-5/0/0 unit 0 tunnel source 192.18.38.1
admin@ex9208# set interfaces gr-5/0/0 unit 0 tunnel destination 192.18.38.2
admin@ex9208# set interfaces gr-5/0/0 unit 0 family inet address
192.18.44.1/24
```

3. Enable OSPF and enable it on the tunnel interface:

```
admin@ex9208# set protocols ospf area 0.0.0.0 interface gr-5/0/0.0
admin@ex9208# commit
```

Cisco commands

1. Ensure the **tunnel** feature is enabled. This step is required for Cisco Nexus 7000 switches, but is not required for Cisco Catalyst switches:

```
Nexus7010# configure terminal
Nexus7010(config)# feature tunnel
```

2. Define IP addresses on the inter-switch link and the link to the Spirent TestCenter traffic generator/analyzer. Note that these addresses are associated with the physical interface, not the GRE tunnel:

```
Nexus7010(config)# interface Ethernet3/9
Nexus7010(config-if)# description physical interface to EX9208 xe-5/0/5
Nexus7010(config-if)# ip address 192.18.38.2/24
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# interface Ethernet3/13
Nexus7010(config-if)# description physical interface to STC
Nexus7010(config-if)# ip address 192.18.98.1/24
Nexus7010(config-if)# no shutdown
```

Note that Cisco Nexus ports are in **shutdown** mode by default, and must be explicitly enabled. This step is not required for Cisco Catalyst switches.

3. Define a GRE tunnel, in this case called **Tunnel0**. In NX-OS, the required commands include a tunnel source, destination, and endpoint IP address:

```
Nexus7010(config-if)# interface Tunnel0
Nexus7010(config-if)# ip address 192.18.44.2/24
Nexus7010(config-if)# tunnel source 192.18.38.2
Nexus7010(config-if)# tunnel destination 192.18.38.1
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# exit
```

3. Enable OSPF and enable it on the tunnel interface's network:

```
Nexus7010(config)# feature ospf
Nexus7010(config)# router ospf 1
Nexus7010(config-rtr)# network 192.18.44.0/24 area 0.0.0.0
Nexus7010(config-rtr)# log-adjacency-changes
Nexus7010(config-rtr)# end
```

Validation

The command "**show interfaces gr-5/0/0**" will verify that the GRE tunnel is up. If run while the Juniper and Cisco devices are forwarding traffic, this command's output will include traffic statistics:

```
admin@ex9200-b> show interfaces gr-5/0/0
Physical interface: gr-5/0/0, Enabled, Physical link is Up
  Interface index: 833, SNMP ifIndex: 1665
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 100000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface gr-5/0/0.0 (Index 818) (SNMP ifIndex 1673)
    Flags: Up Point-To-Point SNMP-Traps 0x4000
    IP-Header 192.18.38.2:192.18.38.1:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    Copy-tos-to-outer-ip-header: Off
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 9154
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.18.44/24, Local: 192.18.44.1, Broadcast:
192.18.44.255

gr-5/0/0, Enabled, Link is Up
Encapsulation: GRE, Speed: 100000mbps
Traffic statistics:                                          Current
delta
  Input bytes:                3281648740 (79102480 bps)
[98256668]
  Output bytes:               2233966800 (39551240 bps)
[49103600]
  Input packets:                 6643094 (20015 pps)
[198901]
  Output packets:                4522200 (10007 pps)
[99400]
```

On Cisco devices, the equivalent command is "**show interface tunnel <tunnel number>**".

Jumbo frame routing

Objective

To validate the ability of Juniper and Cisco switches to correctly route bidirectional traffic with packet lengths greater than 1,500 bytes.

Background

Some routing protocols such as open shortest path first (OSPF) require that both routers agree on the same maximum transmission unit (MTU) when exchanging routing information. For Ethernet interfaces, the requirement for matched MTUs applies equally to jumbo frames (those larger than 1,518 bytes) as to standard-length frames.

In part because of the lack of a standard length for jumbo frames, there is confusion in the marketplace about the maximum frame length possible. Older Linux drivers for Ethernet interfaces in servers support a maximum length of around 7,000 bytes, though most Linux drivers now allow frame lengths of 9,000 bytes or more. Ethernet interfaces of switches and routers typically support a larger protocol data unit (PDU) but there is some confusion as to whether that PDU should be a maximum of 9,000 bytes or 9,216 bytes. Adding to the confusion, implementations differ as to whether the 4-byte cyclic redundancy check (CRC) should or should not be included when stating the maximum frame length.

Juniper and Cisco switches typically support 9,216-byte jumbo frames, including CRC[1]. This section explains how to configure both vendors' devices to set up an OSPF routing session using jumbo frames.

Topology

In this example, a Juniper Virtual Chassis Fabric (comprising two Juniper QFX5100s and one Juniper EX4300) configured as an OSPF router exchanges jumbo frames with a Cisco Nexus 7010 switch. This example uses OSPF because it requires both sides to use the same MTU when exchanging database description messages[2].

Figure 3 illustrates the configuration used to validate jumbo frame routing. In this example, an OSPF routing session will be established between the Juniper and Cisco devices. Both interfaces have IP addresses in the 192.18.64.0/24 subnet.

[1] The Juniper EX9200 supports a maximum Ethernet frame length of 9,192 bytes, or 9,152 bytes when configured in Virtual Chassis mode, to account for a 40-byte internal header. Other Juniper devices, including those given in the examples here, support a maximum Ethernet frame length of 9,216 bytes.
[2] This requirement is specified in RFC 2328, section 10.6.

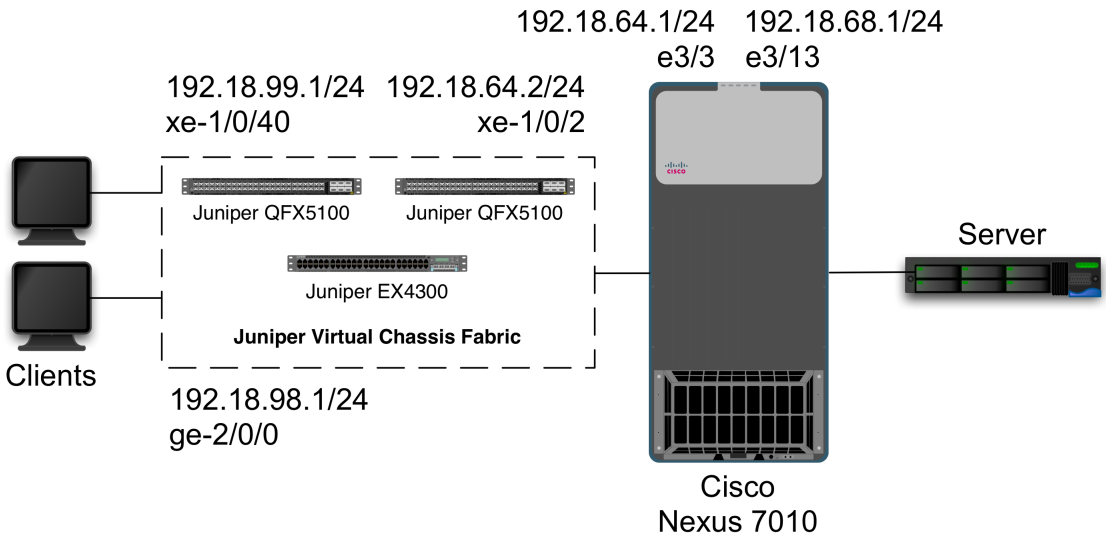

Figure 3: Jumbo frame routing topology

Juniper commands

Jumbo frame support is enabled by adding the **mtu** keyword when configuring interfaces. Note that the **mtu** keyword applies to the physical interface and not the logical unit interface where an IPv4 address is assigned.

Note that the Junos **mtu** keyword does not include the Ethernet CRC. Thus, to pass 9,216-byte Ethernet frames (including CRC), the routing interface will take a command of **mtu 9212**.

These commands assign MTU and IP address to interface xe-0/1/0:

```
admin@VCF> configure
admin@VCF# set interfaces xe-1/0/2 mtu 9212
admin@VCF# set interfaces xe-1/0/2 description "Nexus 7010 e3/3"
admin@VCF# set interfaces xe-1/0/2.0 family inet address 192.18.64.2/24
```

Next, this command starts OSPF routing on interface xe-1/0/2.0. In this example, the interface is a member of OSPF area 0:

```
admin@VCF# set protocols ospf area 0.0.0.0 interface xe-1/0/2.0
admin@VCF# commit
```

Cisco commands

Cisco devices also use the **mtu** keyword in the interface configuration context to enable switching of jumbo frames. Cisco IOS has separate commands for **mtu**, describing the maximum transmission unit for the *Ethernet frame* and for the **ip mtu**, describing the

MTU for the *IP packet.* Cisco NX-OS, as in the Nexus 7010, uses only the **mtu** keyword to cover Ethernet frame length:

```
Nexus7010# configure terminal
Nexus7010(config)# system jumbomtu 9216
Nexus7010(config)# interface Ethernet3/3
Nexus7010(config-if)# description to Juniper VCF xe-1/0/2
Nexus7010(config-if)# mtu 9216
Nexus7010(config-if)# ip address 192.18.64.1/24
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# exit
Nexus7010(config)# router ospf 1
Nexus7010(config-rtr)# log-adjacency-changes
Nexus7010(config-rtr)# network 192.18.64.0 0.0.0.255 area 0
Nexus7010(config-rtr)# end
```

The above example is for Nexus 7000 series switches. On Catalyst 3850 switches, the global configuration **system mtu routing** command sets IP MTU size:

```
Cat3850# configure terminal
Cat3850(config)# system mtu routing 9198
Cat3850(config)# interface TenGigabitEthernet1/0/1
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 10.0.0.1 255.255.255.0
Cat3850(config-if)# exit
Cat3850(config)# router ospf 1
Cat3850(config-rtr)# log-adjacency-changes
Cat3850(config-rtr)# network 10.0.0.0 0.0.0.255 area 0
Cat3850(config-rtr)# end
```

The commands above have been verified with Catalyst 3850 and 3750-E switches routing jumbo frames. On some versions of IOS, the Catalyst 3750 may instead use the global **system mtu jumbo <value>** command.

Validation

Unless both Juniper and Cisco interfaces agree on MTU size, OSPF routing adjacencies will remain in ExStart state, and will never transition to OSPF "full" state. To verify that an OSPF adjacency has entered OSPF "full" state on Juniper switches, use the **show ospf neighbor** command:

```
admin@VCF> show ospf neigbhor
Address          Interface              State     ID               Pri  Dead
192.18.64.1      xe-1/0/2.0             Full      192.18.64.1      1    32
```

On the Cisco device, use the **show ip ospf neighbor** command:

```
Nexus7010# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
192.18.64.2     128   FULL/BDR        00:00:35    192.18.64.2
Ethernet3/3
```

The fact that both routers are in OSPF "Full" state indicates they have agreed to exchange IP packets up to 9,198 bytes long (9,216 bytes, including Ethernet header and CRC). OSPF routing sessions will not be fully established unless both sides agree on MTU size.

Jumbo frame switching

Objective

To validate the ability of Juniper and Cisco switches to correctly switch bidirectional traffic consisting of jumbo frames.

Background

For many years the IEEE Ethernet specification has defined the maximum length of an Ethernet frame to be 1,518 bytes (or 1,522 bytes with an 802.1Q VLAN tag). The use of jumbo frames – those larger than 1,518 bytes – remains nonstandard[3].

In part because of the lack of a standard length for jumbo frames, there is confusion in the marketplace about the maximum frame length possible. Older Linux drivers for Ethernet interfaces in servers support a maximum length of around 7,000 bytes, though most Linux drivers now allow frame lengths of 9,000 bytes or more. Ethernet interfaces of switches and routers typically support a larger protocol data unit (PDU) but there is some confusion as to whether that PDU should be a maximum of 9,000 bytes or 9,216 bytes. Adding to the confusion, implementations differ as to whether the 4-byte cyclic redundancy check (CRC) should or should not be included when stating the maximum frame length.

Juniper and Cisco switches typically support 9,216-byte jumbo frames, including CRC[4]. This section explains how to configure both vendors' switches to exchange jumbo frames.

Topology

In this example, a Juniper Virtual Chassis Fabric switch (comprising two Juniper QFX5100s and one Juniper EX4300) exchanges jumbo frames with a Cisco Nexus 7010. As commonly used in many organizations, VLAN trunk ports connect the switches and VLAN access ports at the edge accept untagged jumbo frames. However, the ability to switch jumbo frames does not depend on VLAN tagging. This example would also work with all interfaces passing untagged traffic.

[3] Recent versions of the 802.3 Ethernet specification have extended the maximum "envelope" frame length to 2,000 bytes to allow for multiple VLAN headers and various encapsulation methods. However, the specification's maximum "basic" frame length remains at 1,518 bytes.

[4] The Juniper EX9200 supports a maximum Ethernet frame length of 9,192 bytes, or 9,152 bytes when configured in Virtual Chassis mode, to account for a 40-byte internal header. Other Juniper devices, including those given in the examples here, support a maximum Ethernet frame length of 9,216 bytes.

Figure 4 illustrates the configuration used to validate jumbo frame switching. As noted in the configuration sections below, all interfaces explicitly support switching of jumbo frames. On the Juniper Virtual Chassis Fabric, two clients (one apiece attached to the Juniper QFX5100 and EX4300) use an untagged VLAN ID of 2001. The 10-Gbit/s Ethernet interface xe-1/0/2 is a trunk port, conveying tagged traffic to the Cisco Nexus 7010 switch. On the Cisco side, interface Ethernet3/3 is also a trunk port. A server is attached to access port Ethernet3/13.

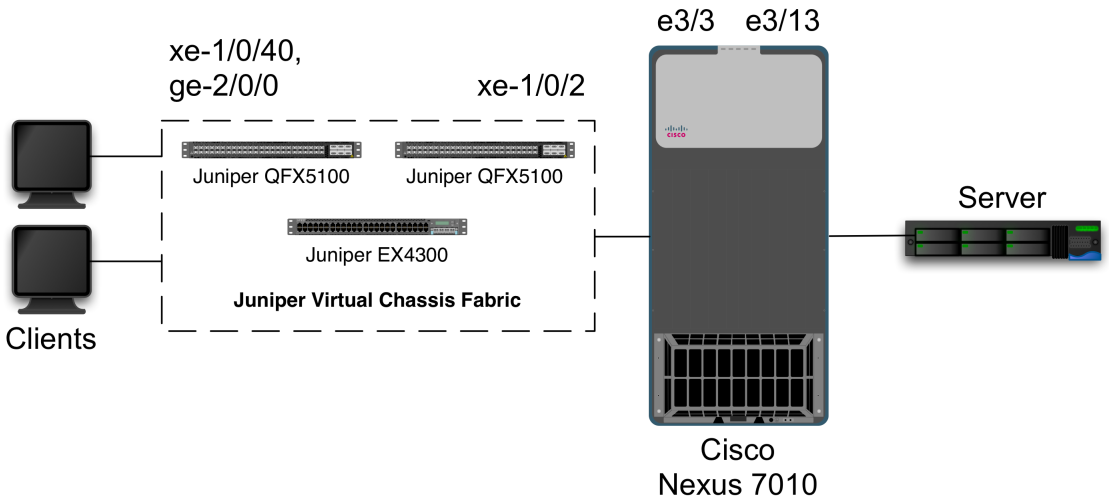

Figure 4: Jumbo frame switching topology

Juniper commands

Jumbo frame support is enabled by adding the **mtu** keyword when configuring interfaces. Note that the **mtu** keyword applies to the physical interface and not the logical unit interface where VLAN membership is assigned.

Note that the Junos **mtu** keyword does not include the Ethernet CRC. Thus, to pass 9,216-byte Ethernet frames (including CRC), untagged (access) ports will take a command of **mtu 9212**, while trunk ports will take a command of **mtu 9216** (to accommodate the 4-byte VLAN tag).

In this example, MTU and VLAN settings are configured separately. First, MTU settings are applied to each interface. Again, note that interface xe-1/0/1 takes a larger MTU value to accommodate VLAN tagging:

```
admin@VCF> configure
admin@VCF# set interfaces xe-1/0/40 mtu 9212
admin@VCF# set interfaces xe-1/0/40 description "to client on QFX5100"
admin@VCF# set interfaces ge-2/0/0 mtu 9212
admin@VCF# set interfaces ge-2/0/0 description "to client on EX4300"
admin@VCF# set interfaces xe-1/0/2 mtu 9216
admin@VCF# set interfaces xe-1/0/2 description "VLAN trunk to Nexus 7010
e3/3"
```

Next, a VLAN is created and interfaces are assigned to the VLAN. In this example, the client-attached interfaces accept untagged traffic, while interface xe-1/0/2 passes tagged traffic to the Cisco switch:

```
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set interfaces xe-1/0/40.0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces ge-2/0/0.0 family ethernet-switching vlan members
v2001
admin@VCF# set interfaces xe-1/0/2.0 family ethernet-switching port-mode
trunk vlan members v2001
```

The spanning tree protocol must be either enabled or disabled on all switches. This command will enable rapid spanning tree on a Juniper switch:

```
admin@VCF# set protocols rstp
admin@VCF# commit
```

To disable rapid spanning tree on a Juniper switch:

```
admin@VCF# set protocols rstp disable
admin@VCF# commit
```

Cisco commands

Cisco devices also use the **mtu** keyword in the interface configuration context to enable switching of jumbo frames. As with the Juniper configuration, VLANs are created separately. Unlike the Juniper example, MTU size and VLAN membership are both associated with the physical interface, and MTU size also is set systemwide. Also, with Cisco IOS and NX-OS devices, the **mtu** keyword does include the Ethernet CRC:

```
Nexus7010# configure terminal
Nexus7010(config)# system jumbomtu 9216
Nexus7010(config)# vlan 2001
Nexus7010(config-vlan)# exit
Nexus7010(config)# interface Ethernet3/3
Nexus7010(config-if)# description VLAN trunk to Juniper VCF
Nexus7010(config-if)# mtu 9216
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# interface Ethernet3/13
Nexus7010(config-if)# description to server
Nexus7010(config-if)# mtu 9216
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode access
Nexus7010(config-if)# switchport access vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

The "**`no shutdown`**" command is mandatory for the Nexus 7010 and other devices running the NX-OS operating system. It is optional for the Cisco Catalyst 3850 and other switches running the IOS operating system.

Validation

Generating a known quantity of jumbo frames between the client and server will validate the ability of the switches to exchange jumbo traffic. This can be verified by examining the interface counters on the switch ports where clients and servers are attached. Alternatively, a test instrument can generate bidirectional jumbo frame traffic between the switches. Both switches should forward all jumbo frames with zero frame loss.

Layer-3 virtual private networks (L3 VPNs)

Objective

To verify the ability of Juniper and Cisco switches to use BGP to create a VPN tunnel over an MPLS-based network.
To verify the ability of Juniper and Cisco switches to forward traffic over an L3 VPN.

Background

MPLS-based VPNs provide virtual, routable IP tunnels across an MPLS network. As described in IETF RFC 4364, customers use BGP routing to set up these tunnels, with no visibility of the service provider's underlying MPLS transport.

MPLS-based VPNs offer advantages over conventional IP-based VPNs for enterprises and service providers in terms of scalability and ease of configuration. Conventional VPNs typically require fully meshed networking among all sites requiring VPN connectivity. Moreover, each router must be reconfigured each time a site is added or deleted. In contrast, a change in one site in an MPLS-based VPN requires reconfiguration only of the service provider's edge router at that site.

Topology

This example models a service provider's MPLS network with provider edge (PE) and provider (P) devices, as well as customer edge (CE) devices representing an enterprise network. In this case, a Juniper MX80 router acts as the P device. A Juniper Virtual Chassis (comprising two Juniper EX9208 switches) and a Cisco Nexus 7010 act as PE devices. A Juniper Virtual Chassis Fabric (comprising two Juniper QFX5100 and one Juniper EX4300 switches) and a Cisco Catalyst 3850 act as CE devices.

The PE devices run MP-BGP (multiprotocol BGP) and OSPF routing protocols, redistributing routes learned from the L3 VPN into OSPF for use by the CE devices. The PE and P devices also run MPLS Label Distribution Protocol (LDP). The PE devices also

run one unique virtual routing and forwarding (VRF) instance for each customer; this allows customers to maintain unique, and possibly overlapping, routing tables (for example, with multiple customers each using the same RFC 1918 private address space). The CE devices run OSPF.

This table lists the IP networks in use in the customer and service provider networks. With L3 VPNs, the customer's equipment has no visibility into the service provider network. Instead, all traffic appears to be routed between customer networks using BGP and OSPF.

Customer networks	Service provider networks
192.18.1.0/24, 192.18.18.0/24, 192.18.40.0/24, 192.18.68.0/24, 192.18.69.0/24, 192.18.98.0/24, 192.18.99.0/24	192.18.38.0/24, 192.38.70.0/24

Figure 5 shows the L3 VPN validation test bed.

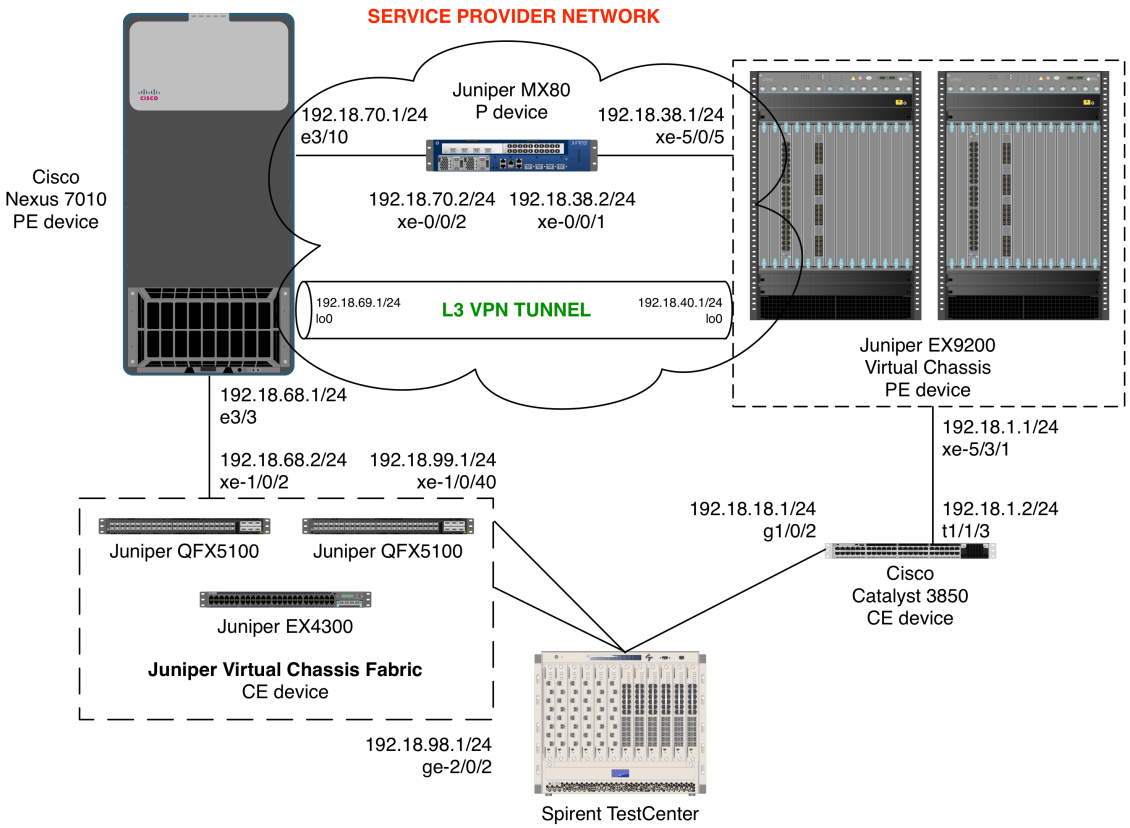

Figure 5: L3 VPN validation topology

Juniper configuration

Juniper Virtual Chassis Fabric (CE device):
With L3 VPNs, no special configuration is needed for CE routers. Indeed, unlike conventional VPNs, the CE routers do not require fully meshed connectivity with all other sites using VPN tunneling.

In this example, the CE devices run OSPF to exchange information with the PE routers. The CE devices also could use static routes or other dynamic routing protocols such as IS-IS or BGP. Again, no MPLS or VPN awareness is required.

1. Configure IP addresses for interfaces xe-1/0/2 (connected with the Cisco Nexus 7000) and xe-1/0/40 and ge-2/0/2 (connected with the Spirent TestCenter test instrument via the Juniper QFX5100 and Juniper EX4300 switches, respectively, within the Virtual Chassis Fabric):

```
VCF> configure
VCF# set interfaces xe-1/0/2 description "to Nexus 7010 int e3/3"
VCF# set interfaces xe-1/0/2 unit 0 family inet address 192.18.68.2/24
VCF# set interfaces xe-1/0/40 description "VCF5100 to stc"
VCF# set interfaces xe-1/0/40 unit 0 family inet address 192.18.99.1/24
VCF# set interfaces ge-2/0/2 description "VCF4300 to stc"
VCF# set interfaces ge-2/0/2 unit 0 family inet address 192.18.98.1/24
```

2. Enable OSPF on the interfaces defined in the previous step:

```
VCF# set protocols ospf area 0.0.0.0 interface xe-1/0/2.0
VCF# set protocols ospf area 0.0.0.0 interface xe-1/0/40.0
VCF# set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
VCF# commit
```

The PE devices will redistribute into BGP any routes learned via OSPF from the CE devices, and the L3 VPN in turn will forward these routes to other sites with VPN tunnels.

On Juniper Virtual Chassis (PE device):
With L3 VPNs, the PE devices require the most extensive configuration. The steps involved include the following:

- Configuring a loopback address
- Configuring addresses on interfaces connected to CE and P devices
- Enabling LDP
- Enabling MPLS
- Configuring a backbone IGP (in this case, OSPF)
- Configuring MP-BGP
- Enabling at least one VRF instance
- Configuring routing protocols for each VRF instance (or static routes)
- Redistributing routes learned from CE devices into MP-BGP

1. Define loopback interface lo0 and configure an IP address for that interface. This address will serve as the PE device's router ID for BGP and LDP:

```
EX9208> configure
EX9208# set interfaces lo0 unit 0 family inet address 192.18.40.1/32
```

2. Define IP addresses for interfaces xe-5/0/5 (connected to the Juniper MX80 acting as a P device) and xe-5/3/1 (connected to the Cisco Catalyst 3850 CE device):

```
EX9208# set interfaces xe-5/0/5 description "to Juniper MX80 P device xe-
0/0/1"
EX9208# set interfaces xe-5/0/5 unit 0 family inet address 192.18.38.1/24
EX9208# set interfaces xe-5/3/1 description "to c3850 CE device int t1/1/3"
EX9208# set interfaces xe-5/3/1 unit 0 family inet address 192.18.1.1/24
```

4. Configure LDP as the MPLS backbone label distribution protocol:

```
EX9208# set protocols ldp interface xe-5/0/5.0
```

5. Enable MPLS on interfaces lo0 and xe-5/0/5:

```
EX9208# set protocols mpls interface xe-5/0/5.0
EX9208# set protocols mpls interface lo0.0
EX9208# set interfaces xe-5/0/5 unit 0 family mpls
```

6. Enable OSPF on the service provider network (IS-IS also would work as an IGP):

```
EX9208# set protocols ospf traffic-engineering
EX9208# set protocols ospf area 0.0.0.0 interface xe-5/0/5.0
EX9208# set protocols ospf area 0.0.0.0 interface lo0.0
```

7. Configure MP-BGP to exchange routes with other PE devices. This example defines a "**Juniper-to-Cisco**" BGP group in which the service provider network uses autonomous system 100 (AS 100):

```
EX9208# set protocols bgp group Juniper-to-Cisco type internal
EX9208# set protocols bgp group Juniper-to-Cisco local-address 192.18.40.1
EX9208# set protocols bgp group Juniper-to-Cisco local-as 100
EX9208# set protocols bgp group Juniper-to-Cisco neighbor 192.18.69.1
family inet-vpn unicast
EX9208# set protocols bgp group Juniper-to-Cisco neighbor 192.18.69.1 peer-
as 100
```

8. Configure a VRF instance. In this example, the VRF's name is "**VPN1**". The "**route-distinguisher**" command uniquely identifies this VRF's network. RDs prevent traffic misrouting when multiple customers use the same network space (for example, when two customers both use net-10 addresses):

```
EX9208# set routing-instances VPN1 instance-type vrf
EX9208# set routing-instances VPN1 route-distinguisher 100:2
EX9208# set routing-instances VPN1 vrf-target target:100:2
EX9208# set routing-instances VPN1 vrf-table-label
```

9. Configure the VRF instance defined in the previous step on interface xe-5/3/1, which connects with the CE device:

```
EX9208# set routing-instances VPN1 interface xe-5/3/1.0
```

10. Configure OSPF for routing between CE and PE devices. This step binds the VRF instance called **VPN1** to the routing protocols or static routes used at customer sites:

```
EX9208# set routing-instances VPN1 protocols ospf area 0.0.0.0 interface
xe-5/3/1.0
EX9208# set routing-instances VPN1 protocols ospf export default-export
EX9208# set routing-instances VPN1 protocols ospf import default-import
```

11. Configure the PE device to redistribute routes learned from CE devices into MP-BGP. This example uses the routes learned from OSPF or BGP, but other routing protocols or static routing could be used with other VRF instances. This example redistributes routes from a policy statement called "**Tsunami1**":

```
EX9208# set policy-options policy-statement Tsunami1-export-policy term 1
from protocol ospf
EX9208# set policy-options policy-statement Tsunami1-export-policy term 1
from protocol bgp
EX9208# set policy-options policy-statement Tsunami1-export-policy term 1
then community add Tsunami1
EX9208# set policy-options policy-statement Tsunami1-export-policy term 1
then accept
EX9208# set policy-options policy-statement Tsunami1-export-policy term 2
then reject
EX9208# set policy-options policy-statement Tsunami1-import-policy term 1
from protocol bgp
EX9208# set policy-options community Tsunami1 members target:100:2
EX9208# commit
```

On Juniper MX80 (P device):

The steps involved for configuration of a P device include the following:

- Configuring a loopback address
- Configuring addresses on interfaces connected to CE and P devices
- Enabling LDP
- Enabling MPLS
- Configuring a backbone IGP (in this case, OSPF)

1. Define loopback interface lo0 and configure an IP address for that interface. This address will serve as the PE device's router ID for BGP and LDP:

```
mx80> configure
mx80# set groups global interfaces lo0 unit 0 family inet address
10.255.3.56/32 primary
```

2. Define IP addresses for interfaces xe-0/0/1 and xe-0/0/2 (connected to the Juniper Virtual Chassis and Cisco Nexus 7010, respectively, each acting as PE devices):

```
mx80# set interfaces xe-0/0/1 description "to Juniper Virtual Chassis xe-
5/0/5"
mx80# set interfaces xe-0/0/1 unit 0 family inet address 192.18.38.2/24
mx80# set interfaces xe-0/0/2 description "to Cisco Nexus 7010 e3/10"
mx80# set interfaces xe-0/0/2 unit 0 family inet address 192.18.70.2/24
```

3. Configure LDP as the MPLS backbone label distribution protocol:

```
mx80# set protocols ldp interface xe-0/0/1.0
mx80# set protocols ldp interface xe-0/0/2.0
mx80# set protocols ldp interface lo0.0
```

4. Enable MPLS on the interfaces connected to PE devices:

```
mx80# set protocols mpls interface xe-0/0/1.0
mx80# set protocols mpls interface xe-0/0/2.0
mx80# set interfaces xe-0/0/1 unit 0 family mpls
mx80# set interfaces xe-0/0/2 unit 0 family mpls
```

5. Enable OSPF on the service provider network. IS-IS would also work as an IGP:

```
mx80# set protocols ospf area 0.0.0.0 interface xe-0/0/1.0
mx80# set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
mx80# set protocols ospf area 0.0.0.0 interface lo0.0
mx80# commit
```

Cisco configuration

Cisco Catalyst 3850 (CE device):

1. Configure IP addresses for interfaces t1/1/3 (connected with the Juniper Virtual Chassis) and g1/0/2 (connected with the Spirent TestCenter test instrument):

```
c3850# configure terminal
c3850(config)# interface TenGigabitEthernet1/1/3
c3850(config-if)# description to Juniper Virtual Chassis int xe-5/3/1
c3850(config-if)# no switchport
c3850(config-if)# ip address 192.18.1.2 255.255.255.0
c3850(config-if)# no shutdown
c3850(config-if)# interface GigabitEthernet1/0/2
c3850(config-if)# description to stc
c3850(config-if)# no switchport
c3850(config-if)# ip address 192.18.18.1 255.255.255.0
c3850(config-if)# no shutdown
c3850(config-if)# exit
```

2. Enable OSPF on the interfaces defined in the previous step:

```
c3850(config)# ip routing
c3850(config)# router ospf 1
c3850(config-rtr)# network 192.18.1.0 0.0.0.255 area 0
```

```
c3850(config-rtr)# network 192.18.18.0 0.0.0.255 area 0
c3850(config-rtr)# log-adjacency-changes
c3850(config-rtr)# end
```

Cisco Nexus 7010 (PE device):

1. Enable the various NX-OS feature sets required for L3 VPNs:

```
Nexus7010# configure terminal
Nexus7010(config)# feature ospf
Nexus7010(config)# feature bgp
Nexus7010(config)# feature mpls l3vpn
Nexus7010(config)# feature mpls ldp
```

2. Define a loopback interface and configure an IP address for that interface. This address will serve as the PE device's router ID for BGP and LDP:

```
Nexus7010(config)# interface loopback0
Nexus7010(config-int)# ip address 192.18.69.1/32
Nexus7010(config-int)# no shutdown
```

Note that an explicit **`no shutdown`** command is mandatory for Cisco Nexus 7000 devices.

3. Define IP addresses for interfaces e3/10 (connected to the Juniper MX80 acting as a P device) and e3/3 (connected to the Juniper Virtual Chassis Fabric):

```
Nexus7010(config-int)# interface e3/10
Nexus7010(config-int)# description to Juniper MX80 P device int xe-0/0/2
Nexus7010(config-int)# ip address 192.18.70.1/24
Nexus7010(config-int)# no shutdown
Nexus7010(config-int)# interface e3/3
Nexus7010(config-int)# description to Juniper Virtual Chassis Fabric
Nexus7010(config-int)# ip address 192.18.68.1/24
Nexus7010(config-int)# no shutdown
Nexus7010(config-int)# exit
```

4. Configure LDP as the MPLS backbone label distribution protocol:

```
Nexus7010(config)# mpls ldp configuration
Nexus7010(config-mpls)# router-id Eth3/10
Nexus7010(config-mpls)# neighbor 192.18.38.1 targeted
Nexus7010(config-mpls)# exit
```

5. Enable MPLS on interface e3/10 (connected to the P device):

```
Nexus7010(config)# interface e3/10
Nexus7010(config-int)# mpls ip
Nexus7010(config-int)# exit
```

6. Enable OSPF on the service provider network (IS-IS also would work as an IGP):

```
Nexus7010(config)# router ospf 1
Nexus7010(config-rtr)# router-id 192.18.69.1
```

```
Nexus7010(config-rtr)# network 192.18.69.1/32 area 0.0.0.0
Nexus7010(config-rtr)# network 192.18.70.0/24 area 0.0.0.0
Nexus7010(config-rtr)# mpls ldp autoconfig area 0.0.0.0
Nexus7010(config-rtr)# mpls ldp sync
Nexus7010(config-rtr)# exit
```

7. Enable BGP and configure MP-BGP to exchange routes with other PE devices. In this example, the service provider network uses autonomous system 100 (AS 100):

```
Nexus7010(config)# router bgp 100
Nexus7010(config-rtr)# router-id 192.18.69.1
Nexus7010(config-rtr)# neighbor 192.18.40.1 remote-as 100
Nexus7010(config-rtr)# address-family vpnv4 unicast
Nexus7010(config-rtr)# send-community extended
Nexus7010(config-rtr)# exit
```

8. Configure a VRF instance. In this example, the VRF's name is "**VPN1**". The "**rd**" command is a route distinguisher that uniquely identifies this VRF's network. RDs prevent traffic misrouting when multiple customers use the same network space (for example, when two customers both use net-10 addresses):

```
Nexus7010(config)# vrf context VPN1
Nexus7010(config-vrf)# rd 100:2
Nexus7010(config-vrf)# address-family ipv4 unicast
Nexus7010(config-vrf)# route-target import 100:2
Nexus7010(config-vrf)# route-target export 100:2
Nexus7010(config-vrf)# exit
```

9. Configure the VRF instance defined in the previous step on interface e3/3, which connects with the CE device:

```
Nexus7010(config)# int e3/3
Nexus7010(config-int)# vrf member VPN1
Nexus7010(config-int)# exit
```

10. Configure OSPF for routing between CE and PE devices. This step binds the VRF instance to the routing protocols or static routes used at customer sites. This example places the learned routes into a route map called **rmap1**:

```
Nexus7010(config)# route-map rmap1 permit 10
Nexus7010(config)# router ospf 1
Nexus7010(config-rtr)# vrf VPN1
Nexus7010(config-rtr)# network 192.18.68.0/24 area 0.0.0.0
Nexus7010(config-rtr)# redistribute bgp 100 route-map rmap1
Nexus7010(config-rtr)# exit
```

11. Configure the PE device to redistribute routes learned from CE devices into MP-BGP. This example uses the routes learned from OSPF process 1, but other routing protocols or static routing could be used with other VRF instances. This example redistributes routes from the route map called **rmap1**:

```
Nexus7010(config)# router bgp 100
Nexus7010(config-router)# vrf VPN1
```

```
Nexus7010(config-router-vrf)# address-family ipv4 unicast
Nexus7010(config-router-vrf)# redistribute direct route-map rmap1
Nexus7010(config-router-vrf)# redistribute ospf 1 route-map rmap1
Nexus7010(config)# end
```

Validation

The Junos command "**`show route protocol bgp`**" will display routes learned via MP-BGP across the service provider's network. The equivalent command for the Cisco Nexus 7010 is "**`show ip route bgp`**". On both devices, the existence of routes learned across the service provider's MPLS network validates that BGP is working in the L3 VPN.

On the data plane, a traffic generator such as Spirent TestCenter should be able to reach sites across the service provider's network across the L3 VPN tunnel.

Link aggregation

Objective

To validate the ability of Juniper and Cisco switches to correctly forward traffic over a logical connection created using IEEE 802.3ad link aggregation.

To verify the ability of Juniper and Cisco switches to use the link aggregation control protocol (LACP) to dynamically remove a member from a link aggregation group (LAG).

Background

The IEEE 802.3ad link specification defines a standards-based method for aggregating multiple physical Ethernet links into a single logical link. The logical link, known as a link aggregation group (LAG), is comprised of multiple *members* (individual pairs of physical interfaces on each switch). LAGs may be defined statically or dynamically, the latter using the link aggregation control protocol (LACP). With LACP enabled, 802.3ad-compliant switches can dynamically add or remove one or more members to a LAG.

Especially when used with LACP, link aggregation adds redundancy to network connections. Depending on the number of flows and the hashing technique used, link aggregation may also boost available bandwidth.

Link aggregation groups also can be defined across multiple chassis, as discussed in the "Multi-Channel Link Aggregation Group" section in this document.

Topology

In this example, a Cisco Catalyst 3850 switch uses a two-member LAG to exchange traffic with a Juniper Virtual Chassis comprised of two Juniper EX9208 switches. The Virtual Chassis could also be a Juniper EX4300, a Juniper QFX510, or both of these

combined into a Virtual Chassis Fabric; these all use the same link aggregation commands given here. Note that LAG member interfaces can reside on different physical switches in a Virtual Chassis or Virtual Chassis Fabric configuration; see the Junos Software Configuration Guide for more details.

Interfaces xe-5/3/1 and xe-12/3/0 on the Juniper switch make up the members of the LAG. On the Cisco switch, the LAG members are interfaces TenGigabitEthernet1/1/3 and TenGigabitEthernet1/1/4. LACP is enabled on all LAG members.

Figure 6 shows the topology used to validate link aggregation and LACP functionality.

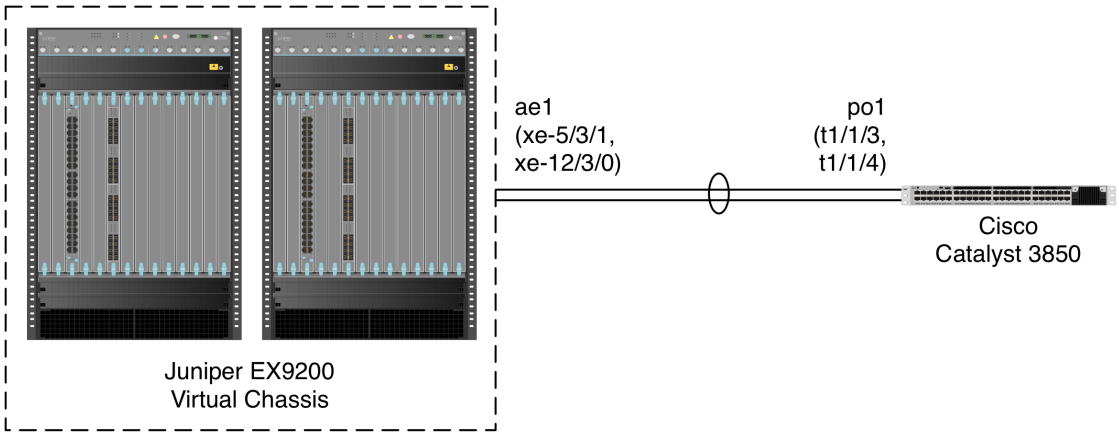

Figure 6: Link aggregation validation topology

Juniper commands

Juniper Junos uses the **ae** interface notation to define each "aggregated Ethernet" instance. The procedure is as follows:

1. Configure the desired number of link aggregation instances (just one, in this example):

```
admin@EX9208> configure
admin@EX9208# set chassis aggregated-devices ethernet device-count 1
```

2. Specify the members to be included within the aggregated Ethernet bundle:

```
admin@EX9208# set interfaces xe-5/3/1 ether-options 802.3ad ae1
admin@EX9208# set interfaces xe-12/3/0 ether-options 802.3ad ae1
```

3. Enable LACP on the aggregated Ethernet instance:

```
admin@EX9208# set interfaces ae1 aggregated-ether-options lacp active
admin@EX9208# set interfaces ae1 description "linkagg to 3850"
```

5. (Optional) Assign the link aggregation interface to be a member of a VLAN. The following example assigns interface **ae1.0** to access mode and allows traffic for VLAN **v2001**:

```
admin@EX9208# set interfaces ae1.0 family ethernet-switching interface-mode
trunk
admin@EX9208# set interfaces ae1.0 family ethernet-switching vlan members
v2001
```

Link aggregation interfaces also can be configured in VLAN trunking mode to carry tagged traffic from multiple VLANs. The following example assigns interface **ae1.0** to trunk-mode membership to carry traffic from VLANs **v2001** and **v2002**:

```
admin@EX9208# set interfaces ae1.0 family ethernet-switching interface-mode
trunk
admin@EX9208# set interfaces ae1.0 family ethernet-switching vlan members
v2001
admin@EX9208# set interfaces ae1.0 family ethernet-switching vlan members
v2002
```

6. To disable or re-enable a member of the LAG, disable that member:

```
admin@EX9208# set interfaces xe-5/3/1 disable
```

Delete the `disable` command to re-enable the LAG member:

```
admin@EX9208# delete interfaces xe-0/1/1 disable
admin@EX9208# commit
```

Cisco commands

1. Create the link aggregation group (called a **port-channel** in Cisco IOS terminology):

```
Cat3850# configure terminal
Cat3850(config)# interface Port-channel1
Cat3850(config-if)# switchport mode access
```

2. Add interfaces to the link aggregation group. The command "**channel-group 1**" adds an interface to the link aggregation group defined in the previous step, while "**mode active**" enables LACP (or "**mode passive**" if the other side of the LAG uses active mode):

```
Cat3850(config)# interface TenGigabitEthernet1/1/3
Cat3850(config-if)# switchport mode access
Cat3850(config-if)# channel-group 1 mode passive
Cat3850(config)# interface TenGigabitEthernet1/1/4
Cat3850(config-if)# switchport mode access
Cat3850(config-if)# channel-group 1 mode passive
Cat3850(config-if)# end
```

These commands are for a Cisco Catalyst 3850. On Cisco Nexus devices running NX-OS, the main difference is that each physical interface must be explicitly enabled with a "**no shutdown**" command. For example, these commands put interface Ethernet3/9 into Port-channel 2:

```
Nexus7010(config)# interface Ethernet3/9
Nexus7010(config-if)# description linkagg to ex9200 xe-5/0/5 and xe-12/0/5
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# channel-group 2 mode passive
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Validation

The Junos command "**show lacp interfaces <aggregated Ethernet interface>**" will show LAG state. The following command was run after disabling interface xe-12/3/0, and validates that LACP on both switches dynamically removed the second member of the LAG. Note that interface xe-12/3/0 is in "Detached" state:

```
admin@EX9200# run show lacp interfaces ae1
Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-5/3/1        Actor    No    No   Yes  Yes  Yes   Yes     Fast
Active
      xe-12/3/0       Partner    No    No   Yes  Yes  Yes   Yes     Slow
Active
      xe-5/3/1        Actor    No   Yes    No   No   No   Yes     Fast
Active
      xe-12/3/0       Partner    No   Yes    No   No   No   Yes     Fast
Passive
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-5/3/1                  Current   Slow periodic Collecting distributing
      xe-12/3/0            Port disabled     No periodic           Detached
```

Link-Layer Discovery Protocol (LLDP)

Objective

To verify the ability of Juniper and Cisco switches to exchange capabilities information using LLDP.

Background

LLDP, as described in the IEEE 802.1AB specification, is a standards-based method of exchanging device capabilities. Unlike Cisco Discovery Protocol (CDP), covered elsewhere in this document, LLDP is an open standard, and thus allows multiple vendors' devices to exchange capabilities data.

Topology

In this example, a Juniper Virtual Chassis comprising two Juniper EX9208 switches uses LLDP to learn the MAC address (chassis ID), port information, and system name of a Cisco Catalyst 3850.

Figure 7 shows the LLDP validation topology. A Juniper Virtual Chassis (comprised of two Juniper EX9208 switches) connects to a Cisco Catalyst 3850 via two 10-gigabit Ethernet interfaces. In this example, each switch is connected with two ports using a link aggregation group (LAG), The LAG also serves as a VLAN trunk port, with traffic using VLAN IDs 2001-2003 allowed. LLDP would also work without link aggregation, and with the two switch ports configured in access mode. Also note that this example assumes spanning tree protocol (STP) has been disabled on both switches, although LLDP would also work with STP enabled.

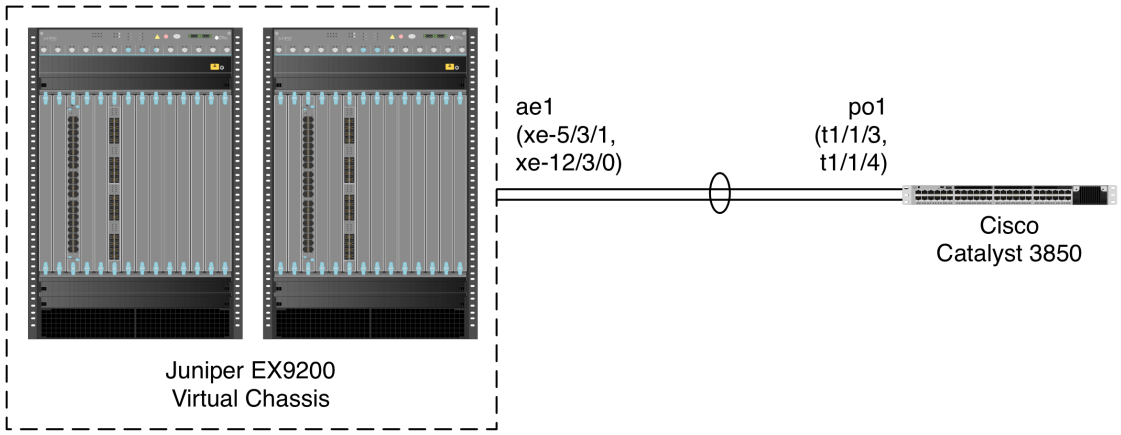

Figure 7: LLDP validation topology

Juniper commands

1. Define VLANs **v2001**, **v2002**, and **v2003**:

```
admin@EX9208> configure
admin@EX9208# set vlans v2001 vlan-id 2001
admin@EX9208# set vlans v2002 vlan-id 2002
admin@EX9208# set vlans v2003 vlan-id 2003
```

2. (Optional) Place interfaces xe-5/1/0 and xe-12/3/0 into trunk mode and allow tagged traffic for the VLANs defined in the previous step. VLAN trunking is optional, and is not required for LLDP to work:

```
admin@EX9208# set interfaces xe-5/3/1 unit 0 family ethernet-switching
interface-mode trunk
admin@EX9208# set interfaces xe-5/3/1 unit 0 family ethernet-switching vlan
members v2001
```

```
admin@EX9208# set interfaces xe-5/3/1 unit 0 family ethernet-switching vlan
members v2002
admin@EX9208# set interfaces xe-5/3/1 unit 0 family ethernet-switching vlan
members v2003
admin@EX9208# set interfaces xe-12/3/0 unit 0 family ethernet-switching
interface-mode trunk
admin@EX9208# set interfaces xe-12/3/0 unit 0 family ethernet-switching
vlan members v2001
admin@EX9208# set interfaces xe-12/3/0 unit 0 family ethernet-switching
vlan members v2002
admin@EX9208# set interfaces xe-12/3/0 unit 0 family ethernet-switching
vlan members v2003
```

3. (Optional) Create link aggregation group (LAG) **ae1** and add interfaces to the LAG. Link aggregation is optional, and is not required for LLDP to work:

```
admin@EX9208# set protocols lldp interface all
admin@EX9208# set interfaces ae1 description "linkagg to 3850"
admin@EX9208# set interfaces ae1 aggregated-ether-options lacp active
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching interface-
mode trunk
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching vlan members
v2001
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching vlan members
v2002
admin@EX9208# set interfaces ae1 unit 0 family ethernet-switching vlan members
v2003
admin@EX9208# set interfaces xe-5/3/1 description "ae1 linkagg to 3850"
admin@EX9208# set interfaces xe-5/3/1 ether-options 802.3ad ae1
admin@EX9208# set interfaces xe-12/3/0 description "ae1 linkagg to 3850"
admin@EX9208# set interfaces xe-12/3/0 ether-options 802.3ad ae1
```

4. (Optional) Enable LLDP. On Juniper switches, LLDP is enabled by default on all interfaces; if LLDP has not been disabled, skip this step. The following (optional) command enables LLDP on all interfaces but it also can be set on a per-interface basis:

```
admin@EX9208# set protocols lldp interface all
```

5. (Optional) Disable rapid spanning tree protocol (RSTP). In this example RSTP is disabled on all interfaces but it also can be set on a per-interface basis:

```
admin@EX9208# delete protocols rstp
admin@EX9208# commit
```

Cisco commands

1. Define VLANs 2001-2003:

```
Cat3850# configure terminal
Cat3850(config)# vlan 2001
Cat3850(config-vlan)# exit
Cat3850(config)# vlan 2002
Cat3850(config-vlan)# exit
Cat3850(config)# vlan 2002
Cat3850(config-vlan)# exit
```

2. (Optional) Define link aggregation group **Port-channel1** and configure it for VLAN trunking. Link aggregation and VLAN trunking are optional, and are not required for LLDP to work:

```
Cat3850(config)# interface Port-channel1
Cat3850(config-if)# description linkagg to EX9208
Cat3850(config-if)# switchport trunk allowed vlan 2001-2003
Cat3850(config-if)# switchport mode trunk
```

3. Define interfaces TenGigabitEthernet1/1/3 and TenGigabitEthernet1/1/4 as VLAN trunk ports and add them to **Port-channel1**:

```
Cat3850(config)# interface range TenGigabitEthernet1/1/3-4
Cat3850(config-if-range)# description linkagg to EX9208
Cat3850(config-if-range)# switchport trunk allowed vlan 2001-2003
Cat3850(config-if-range)# switchport mode trunk
Cat3850(config-if-range)# channel-group 1 mode passive
Cat3850(config-if-range)# exit
```

4. Enable LLDP. On the Cisco Catalyst 3850, this command applies systemwide:

```
Cat3850(config)# lldp run
```

For a Cisco Nexus 7010, the LLDP feature must be enabled. It too applies systemwide:

```
Nexus7010(config)# feature lldp
```

5. Disable spanning tree for VLANs 2001-2003:

```
Cat3850(config)# no spanning-tree vlan 2001-2003
Cat3850(config)# end
```

Validation

On the Juniper switch, the command **`show lldp neighbors`** will verify that the Cisco switch is attached to interfaces xe-5/3/1 and xe-12/3/0:

```
admin@EX9208> show lldp neighbors
Local Interface    Parent Interface    Chassis Id          Port info
System Name
xe-5/3/1           -                   0c:27:24:ce:95:80   Te1/1/3
dc-tme-c3850-01.englab.juniper.net
xe-12/3/0          ae1                 0c:27:24:ce:95:80   Te1/1/4
dc-tme-c3850-01.englab.juniper.net
```

Cisco Catalyst and Nexus switches also use the command **show lldp neighbors:**

```
Cat3850# show lldp neighbors

Capability codes:
    (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
    (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID           Local Intf     Hold-time  Capability      Port ID
EX9208              Te1/1/4        120        B,R             1621
EX9208              Te1/1/3        120        B,R             1591
```

Multi-channel link aggregation group (MC-LAG)

Objective

To verify the ability of multiple Juniper switches to provide multi-channel link aggregation groups, presenting a single logical interface to a Cisco Catalyst access switch and to a Cisco core switch.

Background

To attached switches, an MC-LAG looks and functions the same as standard 802.3ad link aggregation: One or more physical interfaces on the attached switch bond together to form a single logical interface.

The difference with MC-LAG is on the other end: A single logical interface spans multiple physical switches, adding extra resiliency even if an entire switch fails. MC-LAG uses standard LACP messages inside the link aggregation group (LAG), and passes messages between switch chassis to monitor device state.

Switches attached to an MC-LAG require no special configuration beyond the usual link aggregation commands.

Topology

In this example, two Juniper EX9208 chassis form two MC-LAGs, one apiece to a Cisco Catalyst 3850 access switch and to a Cisco Nexus 7010 core switch. Notably, the Juniper core switches do *not* use Juniper Virtual Chassis technology in this example. Instead, each EX9208 is a standalone switch.

Figure 8 illustrates the topology used to validate MC-LAG interoperability.

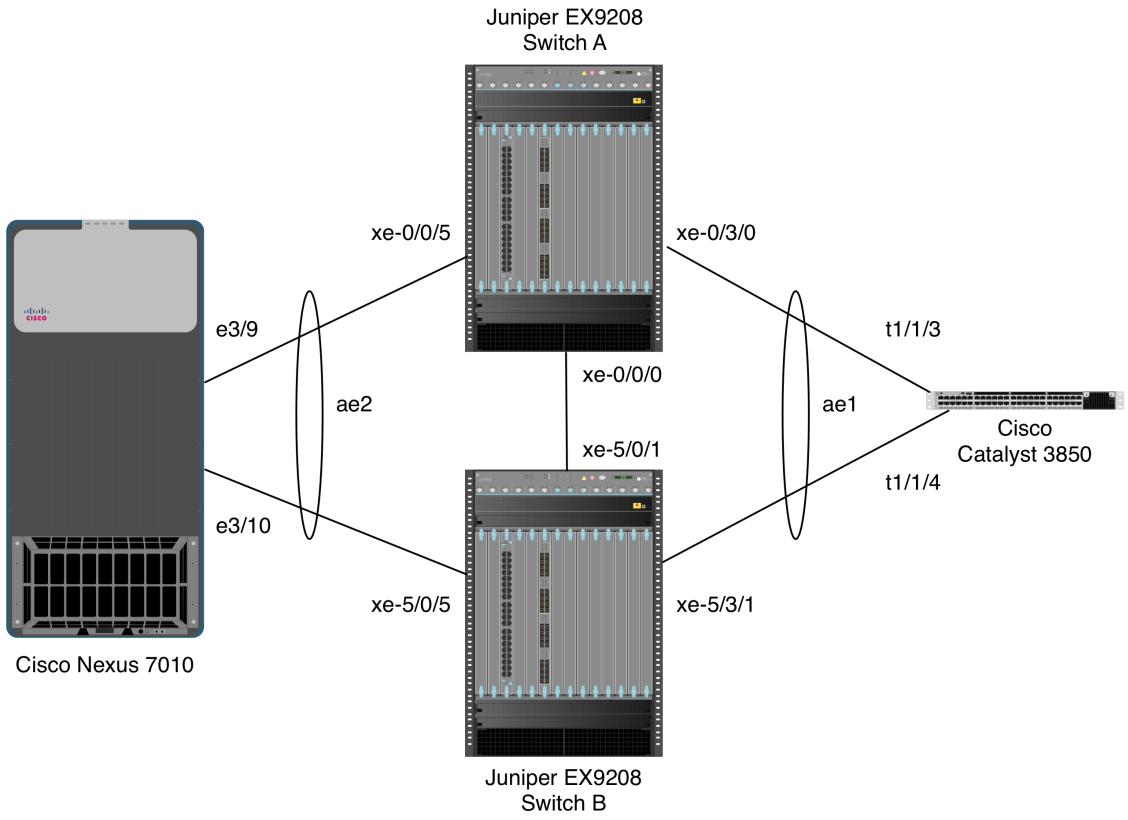

Figure 8: MC-LAG validation topology

Juniper commands

In this example, the steps required are:

- Configure an inter-switch link to carry Inter-Control Center Communications Protocol (ICCP) messages between EX9208 chassis
- Configure an MC-LAG with the Cisco access switch
- Configure an MC-LAG with the Cisco core switch

These steps repeated for Switches A and B.

1. On switch A, begin by configuring an IP address on the loopback interface lo0. This will ensure MC-LAG will continue to function even if the inter-switch link or other EX9208 fails:

```
EX9208A> configure
EX9208A# set interfaces lo0 unit 0 family inet address 192.18.40.2/24
```

2. Next, configure the inter-switch link:

```
EX9208A# set chassis aggregated-devices ethernet device-count 10
EX9208A# set interfaces xe-0/0/0 description "ICCP port to EX9208B"
EX9208A# set interfaces xe-0/0/0 unit 0 family inet address 192.18.39.1/24
EX9208A# set multi-chassis multi-chassis-protection 192.18.39.2 interface xe-
0/0/0
EX9208A# set protocols iccp local-ip-addr 192.18.39.1
EX9208A# set protocols iccp peer 192.18.39.2 session-establishment-hold-time 50
EX9208A# set protocols iccp peer 192.18.39.2 redundancy-group-id-list 1
EX9208A# set protocols iccp peer 192.18.39.2 backup-liveness-detection backup-
peer-ip 192.18.40.1
EX9208A# set protocols iccp peer 192.18.39.2 liveness-detection minimum-
receive-interval 60
EX9208A# set protocols iccp peer 192.18.39.2 liveness-detection transmit-
interval minimum-interval 60
```

3. Configure a link aggregation group **ae1** for connection to the Cisco Catalyst 3850. In this case the LAG is also a trunk that allows traffic for all VLANs, though VLAN configuration is not required for MC-LAG to work:

```
EX9208A# set interfaces xe-0/3/0 description "MC-LAG to 3850"
EX9208A# set interfaces xe-0/3/0 ether-options 802.3ad ae1
EX9208A# set interfaces ae1 aggregated-ether-options lacp active
EX9208A# set interfaces ae1 aggregated-ether-options lacp system-id
00:01:02:03:04:05
EX9208A# set interfaces ae1 aggregated-ether-options lacp admin-key 3
EX9208A# set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
EX9208A# set interfaces ae1 aggregated-ether-options mc-ae redundancy-group
1
EX9208A# set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
EX9208A# set interfaces ae1 aggregated-ether-options mc-ae mode active-
active
EX9208A# set interfaces ae1 aggregated-ether-options mc-ae status-control
active
EX9208A# set interfaces ae1 unit 0 family ethernet-switching interface-mode
trunk
EX9208A# set interfaces ae1 unit 0 family ethernet-switching vlan members
all
```

4. Configure a link aggregation group **ae2** for connection to the Cisco Nexus 7010. In this case the LAG is also a trunk that allows traffic for all VLANs, though VLAN configuration is not required for MC-LAG to work:

```
EX9208A# set interfaces xe-0/0/5 description "MC-LAG to 7010"
EX9208A# set interfaces xe-0/0/5 ether-options 802.3ad ae2
EX9208A# set interfaces ae2 aggregated-ether-options lacp active
EX9208A# set interfaces ae2 aggregated-ether-options lacp system-id
00:01:02:03:04:05
```

```
EX9208A# set interfaces ae2 aggregated-ether-options lacp admin-key 3
EX9208A# set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
EX9208A# set interfaces ae2 aggregated-ether-options mc-ae redundancy-group
1
EX9208A# set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
EX9208A# set interfaces ae2 aggregated-ether-options mc-ae mode active-
active
EX9208A# set interfaces ae2 aggregated-ether-options mc-ae status-control
active
EX9208A# set interfaces ae2 unit 0 family ethernet-switching interface-mode
trunk
EX9208A# set interfaces ae2 unit 0 family ethernet-switching vlan members
all
EX9208A# commit
```

5. On switch B, begin by configuring an IP address on the loopback interface lo0. This will ensure MC-LAG will continue to function even if the inter-switch link or other EX9208 fails:

```
EX9208B> configure
EX9208B# set interfaces lo0 unit 0 family inet address 192.18.40.1/24
```

6. Next, configure the inter-switch link:

```
EX9208B# set chassis aggregated-devices ethernet device-count 10
EX9208B# set interfaces xe-5/0/0 description "ICCP port to EX9208A"
EX9208B# set interfaces xe-5/0/0 unit 0 family inet address 192.18.39.2/24
```

7. Configure a link aggregation group **ae1** for connection to the Cisco Catalyst 3850. In this case the LAG is also a trunk that allows traffic for all VLANs, though VLAN configuration is not required for MC-LAG to work:

```
EX9208B# set interfaces xe-5/3/1 description "linkagg to 3850"
EX9208B# set interfaces xe-5/3/1 ether-options 802.3ad ae1
EX9208B# set interfaces ae1 aggregated-ether-options lacp active
EX9208B# set interfaces ae1 aggregated-ether-options lacp system-id
00:01:02:03:04:05
EX9208B# set interfaces ae1 aggregated-ether-options lacp admin-key 3
EX9208B# set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
EX9208B# set interfaces ae1 aggregated-ether-options mc-ae redundancy-group
1
EX9208B# set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
EX9208B# set interfaces ae1 aggregated-ether-options mc-ae mode active-
active
EX9208B# set interfaces ae1 aggregated-ether-options mc-ae status-control
standby
EX9208B# set interfaces ae1 unit 0 family ethernet-switching interface-mode
trunk
EX9208B# set interfaces ae1 unit 0 family ethernet-switching vlan members
all
```

8. Configure a link aggregation group **ae2** for connection to the Cisco Nexus 7010. In this case the LAG is also a trunk that allows traffic for all VLANs, though VLAN configuration is not required for MC-LAG to work:

```
EX9208B# set interfaces xe-5/0/5 description "linkagg to 7010"
EX9208B# set interfaces xe-5/0/5 ether-options 802.3ad ae2
EX9208B# set interfaces ae2 aggregated-ether-options lacp active
EX9208B# set interfaces ae2 aggregated-ether-options lacp system-id
00:01:02:03:04:05
EX9208B# set interfaces ae2 aggregated-ether-options lacp admin-key 3
EX9208B# set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
EX9208B# set interfaces ae2 aggregated-ether-options mc-ae redundancy-group
1
EX9208B# set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
EX9208B# set interfaces ae2 aggregated-ether-options mc-ae mode active-
active
EX9208B# set interfaces ae2 aggregated-ether-options mc-ae status-control
standby
EX9208B# set interfaces ae2 unit 0 family ethernet-switching interface-mode
trunk
EX9208B# set interfaces ae2 unit 0 family ethernet-switching vlan members
all
EX9208B# commit
```

Cisco commands

As noted, no special MC-LAG configuration is needed on Cisco devices. The only required steps are to define link aggregation groups (called "port channels" in Cisco parlance) on each device, as previously discussed in the "link aggregation" section.

Here are the steps required to configure link aggregation on the Cisco Catalyst 3850:

1. Define a port channel to connect with the MC-LAG. In this example, the interface **Port-channel1** is also a trunk that allows traffic for all VLANs, though VLAN configuration is not required for link aggregation to work:

```
c3850# configure terminal
c3850(config)# interface Port-channel1
c3850(config-if)# switchport mode trunk
c3850(config-if)# exit
```

2. Assign interfaces to **Port-channel1**:

```
c3850(config)# interface TenGigabitEthernet1/1/3
c3850(config-if)# description linkagg to EX9208A
c3850(config-if)# switchport mode trunk
c3850(config-if)# channel-group 1 mode passive
c3850(config-if)# interface TenGigabitEthernet1/1/4
c3850(config-if)# description linkagg to EX9208B
c3850(config-if)# switchport mode trunk
c3850(config-if)# channel-group 1 mode passive
c3850(config-if)# end
```

Here are the steps required to configure link aggregation on the Cisco Nexus 7010:

3. Enable interface configuration and LACP support:

```
Nexus7010# configure terminal
Nexus7010(config)# feature interface-vlan
Nexus7010(config)# feature lacp
```

4. Define a port channel to connect with the MC-LAG. In this example, the interface **Port-channel2** is also a trunk that allows traffic for VLANs 2001-2003, though VLAN configuration is not required for link aggregation to work:

```
Nexus7010(config)# interface Port-channel2
Nexus7010(config-if)# description linkagg to EX9208
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# exit
```

5. Assign interfaces to **Port-channel3**. Note that switches running NX-OS require an explicit "**no shutdown**" command to enable an interface:

```
Nexus7010(config)# interface Ethernet3/9
Nexus7010(config-if)# description linkagg to EX9208A
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# channel-group 2 mode passive
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# interface Ethernet3/10
Nexus7010(config-if)# description linkagg to EX9208B
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# channel-group 2 mode active
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Validation

On the Juniper EX9208 switches, the command "**show interfaces mc-ae**" will display the status of configured MC-LAGs. For example, here is the output of that command on EX9208A:

```
root@EX9200A# run show interfaces mc-ae
 Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.18.39.2 xe-0/0/0.0 up

 Member Link                  : ae2
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae2.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.18.39.2 xe-0/0/0.0 up
```

Similarly, the command "**show iccp**" will display the status of the inter-switch link:

```
root@EX9200A# run show iccp

Redundancy Group Information for peer 192.18.39.2
  TCP Connection       : Established
  Liveliness Detection : Up

Backup liveness peer status: Up
  Redundancy Group ID          Status
    1                           Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1
```

Also, as with regular link aggregation, the command "**show interfaces <aeX> detail**" will display information on the status of LAG *X*, where *X* is the LAG ID.

In the event of a link failure between Juniper and Cisco switches, these same commands will also indicate a change in LAG status.

On Cisco Catalyst 3850 switches , the command "**show etherchannel summary**" will display the status of LAG members. On Cisco Nexus 7010 switches, the equivalent

command is "**`show port-channel summary`**". Again, no MC-LAG awareness or configuration is required on either Cisco switch.

Multicast routing

Objective

To verify the ability of a network comprised of Juniper and Cisco devices to learn multicast routing information using the PIM-SM protocol.

To verify the ability of a network comprised of Juniper and Cisco devices to correctly forward multicast traffic based on routing information learned via PIM-SM.

Background

Protocol Independent Multicast-Sparse Mode (PIM-SM) is a popular choice for multicast routing. Devices running PIM-SM can learn topology information from other PIM-SM routers and make forwarding decisions based on that information.

Like all multicast protocols, PIM-SM uses reverse path forwarding (RPF) lookups to determine which router interface is closest to the multicast source. Because PIM-SM does not include a mechanism to populate an RPF table, it relies on a unicast routing protocol such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) for this purpose.

Topology

This example is similar to that used in the "IP multicast switching" section, with one important change: Here, both Juniper and Cisco devices act as routers running PIM-SM and OSPF. There are no switches, and therefore no IGMP snooping, in this scenario.

In this example, a video server generates multicast traffic to one subnet of a Juniper Virtual Chassis (comprised of two Juniper EX9208s). The Juniper device uses PIM-SM to propagate routing information about that network to other networks, including one in which a Cisco Catalyst 3850 switch, also running PIM-SM, is attached.

Both the Juniper and Cisco devices use PIM-SM and OSPF to propagate routing information. Multicast subscribers attached to routed interfaces, each in a different IP subnet, receive traffic from the streaming video server. The subscriber interfaces also use IGMP (not IGMP snooping) to build a multicast forwarding table.

Figure 9 illustrates the topology used to validate IP multicast routing functionality. PIM-SM and OSPF routing is enabled on both Juniper and Cisco devices.

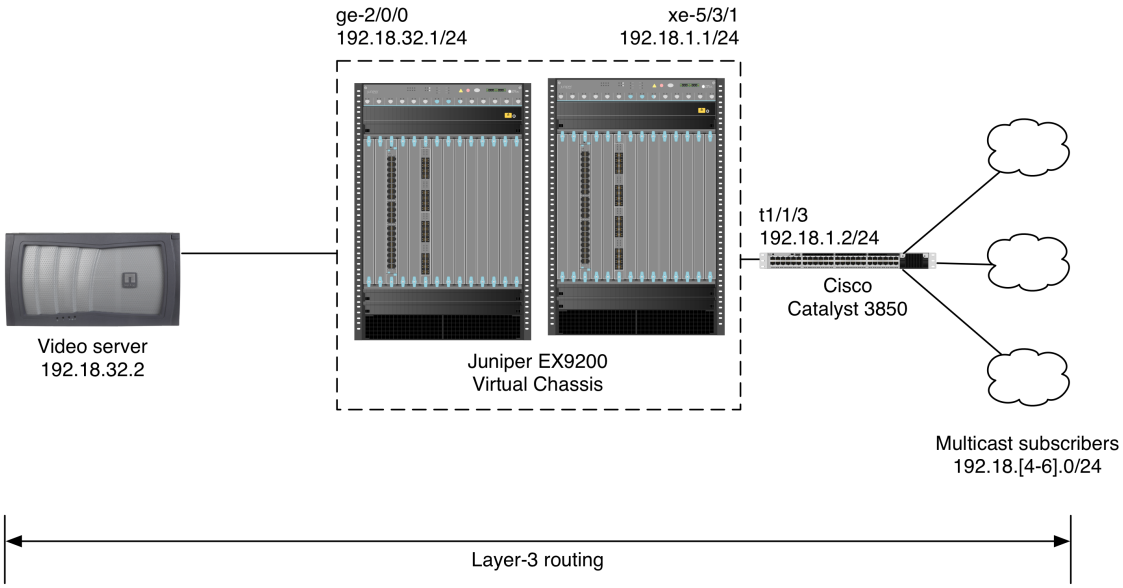

Figure 9: Multicast routing validation topology

Juniper commands

1. Assign IP addresses to the interfaces:

```
admin@VC> configure
admin@VC# set interfaces ge-2/0/0 description "to video server"
admin@VC# set interfaces ge-2/0/0 unit 0 family inet address 192.18.32.1/24
admin@VC# set interfaces xe-5/3/1 description "to Cisco Catalyst 3850"
admin@VC# set interfaces ge-5/3/1 unit 0 family inet address 192.18.1.1/24
```

If the interfaces previously used the **ethernet-switching** keyword, it should be deleted first with the "**delete interfaces <name> family ethernet-switching**" command.

2. Enable PIM-SM on both interfaces:

```
admin@VC# set protocols pim interface ge-2/0/0.0 mode sparse
admin@VC# set protocols pim interface ge-5/3/1.0 mode sparse
```

3. Enable OSPF on both interfaces. This step is not strictly necessary for IP multicast forwarding, but it is required for PIM-SM routing to build an RPF table:

```
admin@VC# set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
admin@VC# set protocols ospf area 0.0.0.0 interface xe-5/3/1.0
```

4. Configure the Juniper device to act as a rendezvous point (RP), in this case by statically assigning an IP address:

```
admin@VC# set protocols pim rp local address 192.18.32.1
admin@VC# commit
```

There are also dynamic methods of discovering and assigning RPs, but those are beyond the scope of this document. *Interdomain Multicast Routing*, mentioned in the "Intended audience" section, covers dynamic RP configuration for Juniper and Cisco devices.

Cisco commands

The following commands apply to a Cisco Catalyst 3850. Except where noted, the syntax is similar for NX-OS devices such as the Cisco Nexus 7010.

1. Enable IP routing and IP multicast routing:

```
Cat3850# configure terminal
Cat3850(config)# ip routing
Cat3850(config)# ip multicast-routing
```

On the Cisco Nexus 7010, the commands are different:

```
Nexus7010# configure terminal
Nexus7010(config)# feature ospf
Nexus7010(config)# feature pim
```

2. Configure IP addresses and PIM-SM for each interface in use:

```
Cat3850(config)# interface TenGigabitEthernet1/1/3
Cat3850(config-if)# description to EX9200 Virtual Chassis
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 192.18.1.2 255.255.255.0
Cat3850(config-if)# ip pim sparse-mode
Cat3850(config-if)# ip igmp version 3
Cat3850(config-if)# no shutdown
Cat3850(config-if)# interface GigabitEthernet1/0/1
Cat3850(config-if)# description to EX9200 Virtual Chassis
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 192.18.4.1 255.255.255.0
Cat3850(config-if)# ip pim sparse-mode
Cat3850(config-if)# ip igmp version 3
Cat3850(config-if)# no shutdown
Cat3850(config-if)# interface GigabitEthernet1/0/2
Cat3850(config-if)# description to EX9200 Virtual Chassis
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 192.18.5.1 255.255.255.0
Cat3850(config-if)# ip pim sparse-mode
Cat3850(config-if)# ip igmp version 3
Cat3850(config-if)# no shutdown
Cat3850(config-if)# interface GigabitEthernet1/0/3
Cat3850(config-if)# description to EX9200 Virtual Chassis
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 192.18.6.1 255.255.255.0
Cat3850(config-if)# ip pim sparse-mode
Cat3850(config-if)# ip igmp version 3
Cat3850(config-if)# no shutdown
Cat3850(config-if)# exit
```

The "**no shutdown**" command is optional with the Catalyst 3850, but mandatory for the Nexus 7010 and other devices running the NX-OS operating system.

Also, note that in this example IP addresses are configured directly on physical interfaces. If desired, IP addresses and routing information can instead be assigned to VLAN interfaces. In that case, each physical interface should be put into **switchport** mode and assigned to a VLAN. The "VLAN Trunking" section of this document has more details.

3. Configure OSPF. This step is not strictly necessary for IP multicast forwarding, but it is required for PIM-SM routing to build an RPF table:

```
Cat3850(config)# router ospf 1
Cat3850(config-rtr)# log-adjacency-changes
Cat3850(config-rtr)# network 192.18.1.0 0.0.0.255 area 0
Cat3850(config-rtr)# network 192.18.4.0 0.0.0.255 area 0
Cat3850(config-rtr)# network 192.18.5.0 0.0.0.255 area 0
Cat3850(config-rtr)# network 192.18.6.0 0.0.0.255 area 0
Cat3850(config-if)# exit
```

4. Configure a PIM rendezvous point (RP), which in this case was statically defined on the Juniper device:

```
Cat3850(config)# ip pim rp-address 192.18.32.1
Cat3850(config-if)# end
```

Validation

Once subscribers attached to the Cisco Catalyst 3850 have joined multicast groups by sending IGMPv3 reports with join messages, any multicast traffic for those groups offered to interface xe-1/0/40 on the Juniper Virtual Chassis will be forwarded to all subscriber ports on the Cisco Catalyst 3850.

The Junos command **show pim neighbors brief** also will verify that the Juniper and Cisco devices see one another and can exchange topology updates. On Cisco devices, the equivalent command is "**show ip pim neighbors**".

Multicast switching

Objective

To verify the ability of Juniper and Cisco switches to correctly forward traffic from a network using IGMP snooping for multicast switching.

Background

Ethernet switches use Internet group management protocol (IGMP) snooping to determine where a switch should forward multicast traffic. With IGMP snooping enabled, a switch listens for IGMP reports from attached devices that wish to receive multicast traffic. The switch then maps subscribed multicast group address(es) to the interface on which the subscriber is attached. When the switch receives traffic destined for an IP multicast group address, it will forward it only to those interfaces from which it has heard membership reports.

Topology

In this example, both Juniper and Cisco switches operate purely in Layer-2 mode, with no multicast or unicast routing protocols configured. This test case assumes routing is handled elsewhere in the network.

A video server generates multicast traffic that is routed across a network and reaches the two switches described here. The switches, in turn, use their IGMP snooping tables to determine which ports should and should not receive multicast traffic.

The streaming video server sends traffic to 10 multicast group addresses in the range of 225.0.1.0 through 225.0.1.9. Subscribers attached to the Juniper and Cisco switches join all 10 multicast groups.

Figure 10 illustrates the topology used to validate IP multicast switching functionality. Both the Juniper and Cisco switches use IGMP snooping.

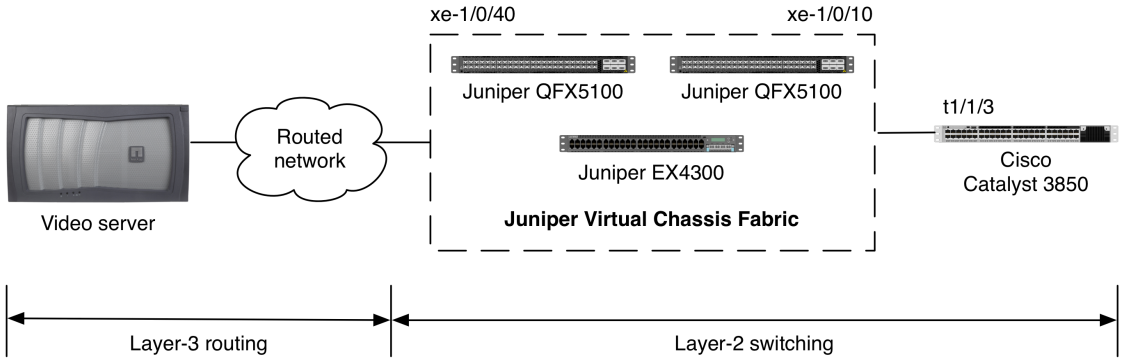

Figure 10: Multicast switching validation topology

Juniper commands

In this example, IGMP snooping is enabled on a per-VLAN basis. For example, these commands create a VLAN called “v2001” and then enable IGMP snooping on that VLAN:

```
admin@VCF> configure
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set protocols igmp-snooping vlan v2001
```

If desired, IGMP snooping can be enabled on all VLANs using the **all** keyword:

```
admin@VCF# set protocols igmp-snooping vlan all
```

Next, configure both Juniper ports to use VLAN trunking and to allow traffic from VLAN "v2001":

```
admin@VCF# set interfaces xe-1/0/10 description "to 3850"
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces xe-1/0/40 description "to router"
admin@VCF# set interfaces xe-1/0/40 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/40 unit 0 family ethernet-switching vlan
members v2001
```

Finally, enable IGMP version 3, the most recent version of the protocol on both interfaces:

```
admin@VCF# set protocols igmp interface xe-1/0/10.0 version 3
admin@VCF# set protocols igmp interface xe-1/0/40.0 version 3
admin@VCF# commit
```

Cisco commands

The following commands apply to a Cisco Catalyst 3850. Except where noted, the syntax is similar for Nexus 7010 switches.

1. Create VLAN 2001:

```
Cat3850# configure terminal
Cat3850(config)# vlan 2001
Cat3850(config-vlan)# exit
```

2. Configure interface TenGigabitEthernet1/1/3 (connected to the Juniper Virtual Chassis Fabric) for VLAN trunking and to allow VLAN 2001:

```
Cat3850(config)# interface TenGigabitEthernet1/1/3
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# switchport trunk allowed vlan 2001
Cat3850(config-if)# no shutdown
Cat3850(config-if)# exit
```

The "**no shutdown**" command is optional with the Catalyst 3850, but mandatory for the Nexus 7010 and other devices running the NX-OS operating system.

3. Set the IP address of the IGMP querier. Usually this is a multicast-enabled router near the switches:

```
Cat3850(config)# ip igmp snooping querier address 192.18.36.1
Cat6500(config)# end
```

Validation

On the Juniper Virtual Chassis Fabric, results of the command **`show igmp-snooping membership`** will verify that the switch has correctly mapped multicast groups to the appropriate subscriber interfaces.

```
root@VCF> show igmp snooping membership | no-more
Instance: default-switch

Vlan: v2001

Learning-Domain: default
Interface: xe-1/0/10.0, Groups: 10
    Group: 225.0.1.0
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 10.205.0.2
        Group timeout:     223 Type: Dynamic
    Group: 225.0.1.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 10.205.0.2
        Group timeout:     223 Type: Dynamic
..
Interface: xe-1/0/40.0, Groups: 0
```

The output for interface xe-1/0/10 will continue through multicast group address 225.0.1.9. Note that interface xe-1/0/40 shows no associated multicast groups. This is because all subscribers are connected to the Cisco switch. There are no multicast group subscribers directly attached to the Juniper switch. Thus, incoming packets with a destination address of an IP multicast group should only be forwarded through interface xe-1/0/10.

The command **`show interface <name> extensive`** will verify correct forwarding of multicast traffic. The "Multicast packets" counter (under "MAC statistics") will increment on interfaces with multicast subscribers attached, and will not increment on other interfaces.

On Cisco devices, the command "**`show ip igmp snooping groups`**" will display information about multicast group membership.

Real-Time Performance Monitoring (RPM)

Objective

To validate the ability of Juniper EX Series switches to perform real-time health checks on attached devices.

Background

Juniper's Real-Time Performance Monitoring (RPM) feature can perform "health checks" on attached network devices and servers using ICMP, HTTP, TCP and UDP probes and requests. These active probes monitor devices across any network and investigate reachability problems. RPM keeps a history of the most recent 50 probes; such monitoring over time can be useful in troubleshooting and capacity planning.

Topology

In this example, a Juniper Virtual Chassis comprising two Juniper EX9208 switches uses uses ICMP probes to monitor round-trip times between it and a Cisco Catalyst 3850E switch. Note that no RPM-specific commands are needed on the Cisco switch. The same RPM configuration on the Juniper switch will work with any Cisco switch, or indeed any remote device capable of responding to pings (ICMP probe requests).

Figure 11 illustrates the RPM validation test bed. The Juniper Virtual Chassis in this example uses an IP address of 192.18.0.10/24 assigned to interface xe-5/3/1. This interface sends ICMP probe requests to a Cisco switch with an address of 192.18.0.1/24. The same RPM configuration would work in a Juniper switch configuration in which a VLAN is created and an IP address is assigned to the VLAN.

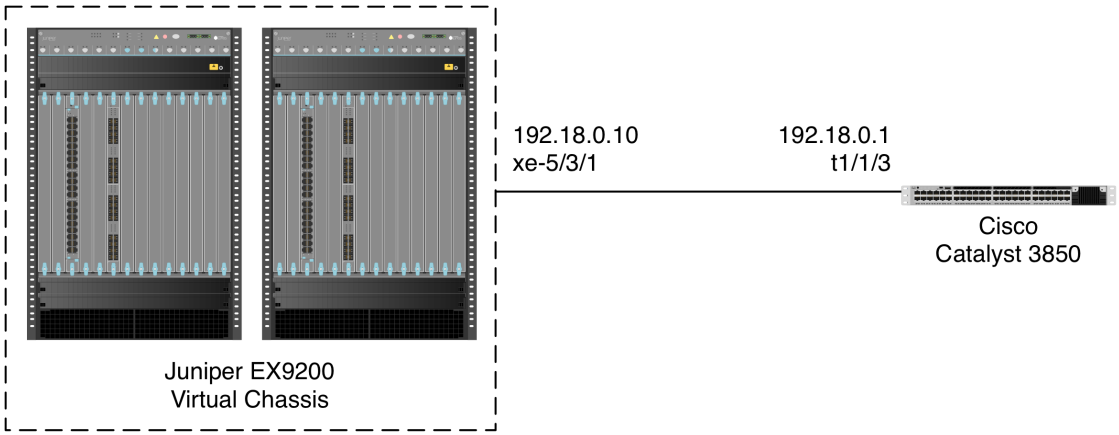

Figure 11: Real-Time Performance Monitoring validation topology

Juniper configuration

1. Assign an IP address of 192.18.0.10/24 to interface xe-5/3/1:

```
admin@EX9208> configure
admin@EX9208# set interfaces ge-0/0/22 unit 0 family inet address
192.18.0.10/24
```

2. Define an RPM probe and test for the Cisco switch at 192.18.0.1:

```
admin@EX9208# set services rpm probe myprobe test t1 probe-type icmp-ping-
timestamp
admin@EX9208# set services rpm probe myprobe test t1 target address
192.18.0.1
admin@EX9208# set services rpm probe myprobe test t1 probe-count 10
admin@EX9208# set services rpm probe myprobe test t1 probe-interval 1
admin@EX9208# set services rpm probe myprobe test t1 test-interval 15
admin@EX9208# set services rpm probe myprobe test t1 data-size 128
admin@EX9208# commit
```

Note that although only ICMP is used here, a single probe can encompass multiple tests using multiple types of health checks.

Cisco configuration

1. Assign an IP address of 192.18.0.1/24 to the monitored interface (in this case TenGigabitEthernet1/1/3):

```
Cat3850# configure terminal
Cat3850(config)# interface GigabitEthernet1/1/3
Cat3850(config-if)# no switchport
Cat3850(config-if)# ip address 192.18.0.1 255.255.255.0
Cat3850(config-if)# end
```

Cisco Nexus 7000 switches running the NX-OS operating system also require an explicit "**`no shutdown`**" command to enable an interface.

No RPM-specific configuration is needed on the Cisco switch, or on any other device monitored using RPM.

Validation

The command **`show services rpm history-results`** will display up to 50 results of RPM probes and tests:

```
admin@ex9208# run show services rpm history-results

    Owner, Test                 Probe received              Round trip time
    myprobe, t1                 Fri Apr 18 00:31:20 2014            2802
usec
    myprobe, t1                 Fri Apr 18 00:31:21 2014            5441
usec
    myprobe, t1                 Fri Apr 18 00:31:22 2014            9909
usec
    myprobe, t1                 Fri Apr 18 00:31:23 2014            2754
usec
    myprobe, t1                 Fri Apr 18 00:31:24 2014            3334
usec
```

Redundant Trunk Group (RTG)

Objective

To validate failover functionality of Juniper's Redundant Trunk Group (RTG) feature between Juniper and Cisco switches.

Background

Juniper's Redundant Trunk Group (RTG) feature allows definition of primary and secondary paths between switches and redirects traffic across the secondary trunk if the primary link fails. RTG provides an alternative to spanning tree bridging for redundancy. RTG works in mixed Juniper-Cisco environments with no additional configuration needed on Cisco switch ports. Up to 16 redundant trunk groups can be defined on a standalone switch or Juniper Virtual Chassis Fabric.

Topology

In this example, three devices – one from Juniper and two Cisco – form a ring topology. The devices are the Juniper Virtual Chassis Fabric (in turn comprised of two Juniper QFX5100 and one Juniper EX4300 switches); a Cisco Nexus 7010; and a Cisco Catalyst 3850. At test time, RTG was not supported in the Juniper EX9200, so it was not included in this test.

In this example, the inter-switch ports are VLAN trunk ports allowing traffic from VLAN IDs 2001-2003. However, RTG works equally well with or without VLAN trunking.

Spirent TestCenter traffic generator/analyzers offer frames to access ports on each switch.

Spanning tree, which is enabled by default on Juniper and Cisco switches, is disabled in this example. Instead RTG configured on the Juniper Virtual Chassis Fabric sets up primary and secondary traffic paths. When trunk links are configured as part of an RTG, they cannot be part of a spanning tree topology.

Figure 12 shows the RTG test bed topology.

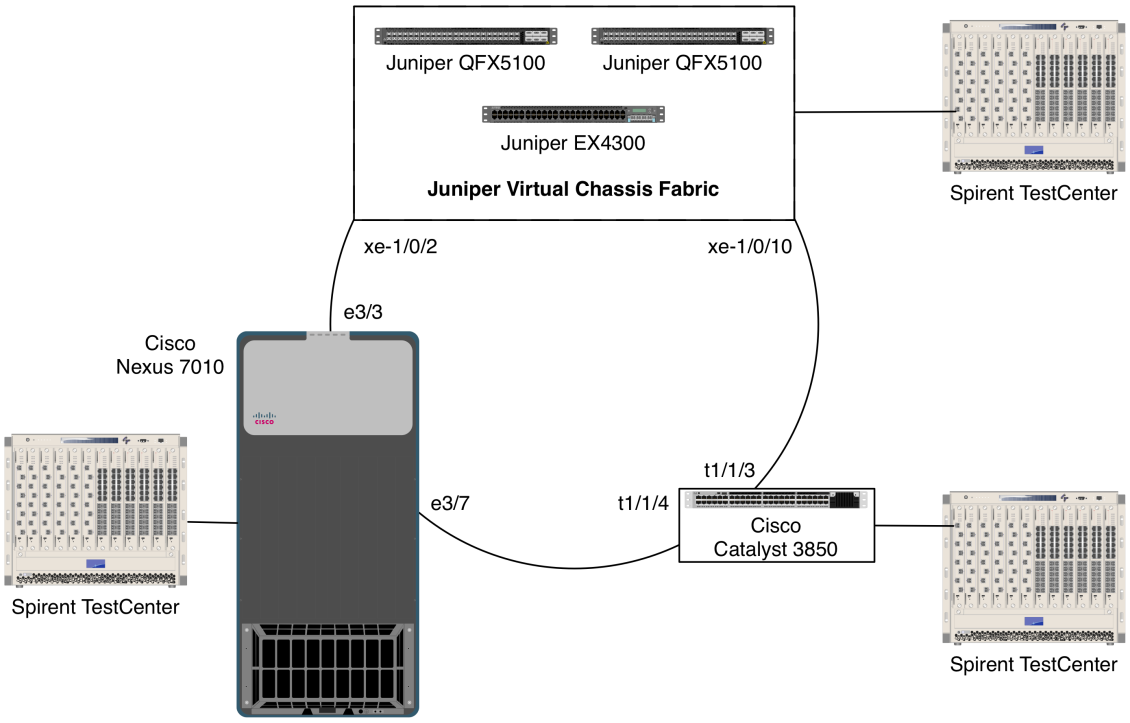

Figure 12: Redundant Trunk Group validation topology

Initially, ports xe-1/0/2 on the Juniper Virtual Chassis Fabric is defined as the primary path for the RTG. If a link failure occurs, the Virtual Chassis Fabric will use the other trunk port on the same switch.

Juniper commands

1. Define three VLANs with VLAN IDs of 2001-2003:

```
admin@VCF> configure
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set vlans v2002 vlan-id 2002
admin@VCF# set vlans v2003 vlan-id 2003
```

2. Define interfaces xe-1/0/2.0 and xe-1/0/10.0 as VLAN trunk ports allowing traffic from the three VLANs defined in the previous step:

```
admin@VCF# set interfaces xe-1/0/2 description "RTG to Nexus 7010 e3/3"
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2002
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2003
admin@VCF# set interfaces xe-1/0/10 description "RTG to Cat3850 t1/1/4"
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching vlan
members v2002
admin@VCF# set interfaces xe-1/0/10 unit 0 family ethernet-switching vlan
members v2003
```

Note that this example uses trunk ports. Access ports also can be members of redundant trunk groups.

3. Disable rapid spanning tree, which is enabled by default:

```
admin@VCF# delete protocols rstp
```

4. Define an RTG named "**rtg0**" and set interface xe-1/0/2.0 as the primary path and interface xe-1/0/10.0 as the secondary path:

```
admin@VCF# set switch-options redundant-trunk-group group rtg0 interface
xe-1/0/2.0 primary
admin@VCF# set switch-options redundant-trunk-group group rtg0 interface
xe-1/0/10.0
admin@VCF# commit
```

Cisco commands

Cisco Nexus 7010:

1. NX-OS switch configuration does not require any RTG-specific commands. Simply define VLANs; disable spanning tree on those VLANs; and assign switch ports to be trunk-mode members of those VLANs:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001-2003
Nexus7010(config-vlan)# exit
Nexus7010# no spanning-tree vlan 2001-2003
Nexus7010# interface Ethernet3/3
Nexus7010(config-if)# description RTG to VCF xe-1/0/2
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# no shutdown
Nexus7010# interface Ethernet3/7
Nexus7010(config-if)# description RTG to Cat3850 t1/1/4
```

```
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Note that Cisco Nexus 7000 switches require an explicit "**no shutdown**" command to enable an interface.

Cisco Catalyst 3850:

1. Cisco Catalyst switches do not require any RTG-specific configuration. Simply define VLANs; disable spanning tree on those VLANs; and assign switch ports to be trunk-mode members of those VLANs:

```
Cat3850# configure terminal
Cat3850(config)# vlan 2001-2003
Cat3850(config-vlan)# exit
Cat3850(config)# no spanning-tree vlan 2001-2003
Cat3850(config)# interface TenGigabitEthernet1/1/3
Cat3850(config-if)# description RTG to Juniper VCF xe-1/0/10
Cat3850(config-if)# switchport
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# switchport trunk allowed vlan 2001-2003
Cat3850(config)# interface TenGigabitEthernet1/1/4
Cat3850(config-if)# description RTG to Nexus 7010 e3/7
Cat3850(config-if)# switchport
Cat3850(config-if)# switchport mode trunk
Cat3850(config-if)# switchport trunk allowed vlan 2001-2003
Cat3850(config-if)# end
```

Validation

The **show redundant-trunk-group** command indicates the current RTG state. This example is from the Juniper Virtual Chassis Fabric:

```
admin@VCF# run show redundant-trunk-group group rtg0

Interface   State       Bandwidth   Time of last flap
Flap

count

xe-1/0/2.0  Up/Pri/Act  10 Gbps     Never
0
xe-1/0/10.0 Up          10 Gbps     2014-04-18 01:38:23  (00:01:55 ago)
6
```

Note that interface xe-1/0/2.0 is the primary path. After offering traffic from Spirent TestCenter, the packet counters for interfaces xe-1/0/2.0 and xe-1/0/10.0 will indicate that the switch forwarded all traffic to interface xe-1/0/2.0, the primary path in the RTG.

To verify correct operation of RTG redundancy, disable the primary path:

```
admin@VCF# set interfaces xe-1/0/2 disable
admin@VCF# commit
```

Now the **show redundant-trunk-group** command will indicate the primary interface is down while the secondary interface remains up:

```
admin@VCF# show redundant-trunk-group
root@VCF# run show redundant-trunk-group group rtg0

Interface   State       Bandwidth   Time of last flap
Flap

count

xe-1/0/2.0  Dwn/Pri     10 Gbps     2014-04-18 01:47:25  (00:00:11 ago)
1
xe-1/0/10.0 Up/Act      10 Gbps     2014-04-18 01:38:23  (00:09:13 ago)
6
```

Note also that the command output indicates when the primary interface went down and that its flap count has incremented by 1.

Again, after offering traffic from Spirent TestCenter, the packet counters for both interfaces will indicate that the switch forwarded all traffic to interface xe-1/0/10.0, the secondary path in the RTG.

Spanning tree case 1: Rapid spanning tree protocol (RSTP)

Objective

To verify interoperability of a rapid spanning tree topology between Juniper and Cisco switches.
To measure convergence time of a rapid spanning tree topology between Juniper and Cisco switches after link failure.

Background

The spanning tree protocol is widely used in Ethernet networks for loop prevention and redundancy. Rapid spanning tree, defined in IEEE 802.1w, provides much faster convergence time after a link or device failure than the original 802.1D spanning tree specification.

Topology

This example uses redundant links between two Juniper switches and one Cisco switch. Junos running on Juniper EX switches supports rapid spanning tree protocol (RSTP) by default. Cisco switches also support spanning tree by default; although Cisco IOS defines

the spanning tree mode as that vendor's proprietary "PVST Plus" mode, it is interoperable with other vendors' rapid spanning tree implementations.

Figure 13 illustrates the RSTP validation test bed. This example involves a standalone Juniper EX4300 switch; a Juniper Virtual Chassis Fabric (comprised of two Juniper QFX5100 and one Juniper EX4300 switches); and a Cisco Nexus 7010.

All inter-switch links use VLAN trunk ports that allow tagged traffic with a VLAN ID of 2001. Rapid spanning is enabled by default on the Juniper switches. VLAN trunking is not required by spanning tree; it would work equally well with interfaces in access mode.

Cisco's "PVST Plus," enabled by default on the Nexus 7010, is interoperable with standard rapid spanning tree. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology.

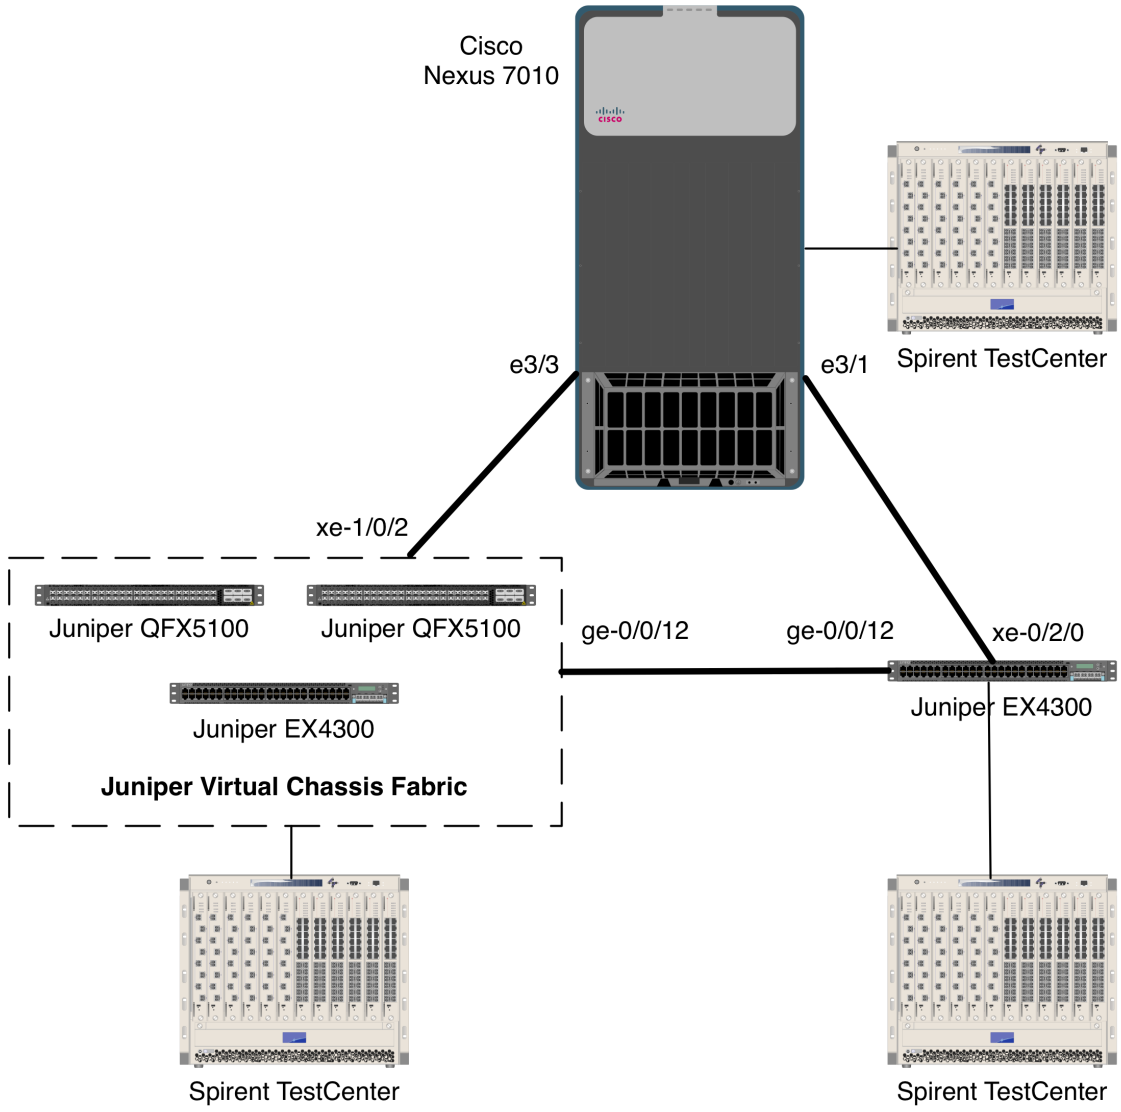

Figure 13: RSTP validation topology

Juniper commands

Juniper Virtual Chassis Fabric:

1. Create a VLAN with an ID of 2001. Then assign interfaces to carry traffic from that VLANs in trunking mode:

```
admin@VCF> config
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set interfaces xe-1/0/2 description "RSTP to Nexus 7010 e3/3"
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces ge-0/0/12 description "RSTP to Juniper EX4300 ge-
0/0/12"
```

```
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
v2002
```

2. Enable rapid spanning tree. On a new switch configuration, this step should be unnecessary since rapid spanning tree is enabled by default:

```
admin@VCF# set protocols rstp
admin@VCF# commit
```

Juniper EX4300:

1. Create a VLAN with an ID of 2001. Then assign interfaces to carry traffic from that VLANs in trunking mode:

```
admin@EX4300> config
admin@EX4300# set vlans v2001 vlan-id 2001
admin@EX4300# set interfaces xe-0/2/0 description "RSTP to Nexus 7010 e3/1"
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members v2001
admin@EX4300# set interfaces ge-0/0/12 description "RSTP to Juniper EX4300
ge-0/0/12"
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members v2001
```

2. Enable rapid spanning tree. On a new switch configuration, this step should be unnecessary since rapid spanning tree is enabled by default:

```
admin@EX4300# set protocols rstp
admin@EX4300# commit
```

Cisco commands

1. Create a VLAN with an ID of 2001. Then assign interfaces to carry traffic from that VLANs in trunking mode:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001
Nexus7010(config-vlan)# exit
Nexus7010(config)# interface Ethernet3/1
Nexus7010(config-if)# description STP to EX4300 xe-0/2/0
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/3
Nexus7010(config-if)# description STP to VCF xe-1/0/2
Nexus7010(config-if)# switchport
```

```
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# exit
```

Note that Cisco Nexus 7000 switches require the explicit "**no shutdown**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

2. Enable PVST Plus. On a new switch configuration, this step should be unnecessary since PVST Plus is enabled by default:

```
Nexus7010(config)# spanning-tree mode rstp
Nexus7010(config)# end
```

On Cisco Catalyst switches, the equivalent command is "**spanning-tree mode rapid-pvst**".

Validation

The command **show spanning-tree bridge brief** will display a summary of spanning tree parameters:

```
admin@VCF# run show spanning-tree bridge brief
STP bridge parameters
Routing instance name               : GLOBAL
Context ID                          : 0
Enabled protocol                    : RSTP
...
```

On the Cisco Nexus 7010, the equivalent command is "**show spanning-tree <vlan ID>**":

```
Nexus7010# show spanning-tree vlan 2001

VLAN2001
  Spanning tree enabled protocol rstp
  Root ID    Priority    34769
             Address     0024.f718.9ec1
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    34769  (priority 32768 sys-id-ext 2001)
             Address     0024.f718.9ec1
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth3/1           Desg FWD 2         128.385  P2p
Eth3/3           Back BLK 2         128.387  P2p
```

Note that interface e3/1 is in forwarding state, and e3/3 is in blocking state.

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then rapid spanning tree convergence time is 47 ms.

Spanning tree case 2: Multiple spanning tree protocol (MSTP)

Objective

To verify interoperability of a multiple spanning tree topology between Juniper and Cisco switches.
To measure convergence time of a multiple spanning tree topology between Juniper and Cisco switches after link failure.

Background

As defined in IEEE specification 802.1s, the multiple spanning tree protocol (MSTP) adds loop prevention and redundancy on a per-VLAN basis. With MSTP, individual spanning tree topologies can be configured for each VLAN.

Topology

This example uses redundant links between two Juniper switches and one Cisco switch. VLAN IDs of 2001 and 2002 have been defined on both the Juniper and Cisco switches, and MSTP is enabled on all switches.

Figure 14 illustrates the MSTP validation test bed. This example involves a standalone Juniper EX4300 switch; a Juniper Virtual Chassis Fabric (comprised of two Juniper QFX5100 and one Juniper EX4300 switches); and a Cisco Nexus 7010.

As shown in the figure, the links interconnecting each switch are trunk ports that allow tagged traffic from VLAN IDs 2001 and 2002. Two access-mode ports are configured on each switch: One apiece for VLAN IDs 2001 and 2002. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology in each VLAN.

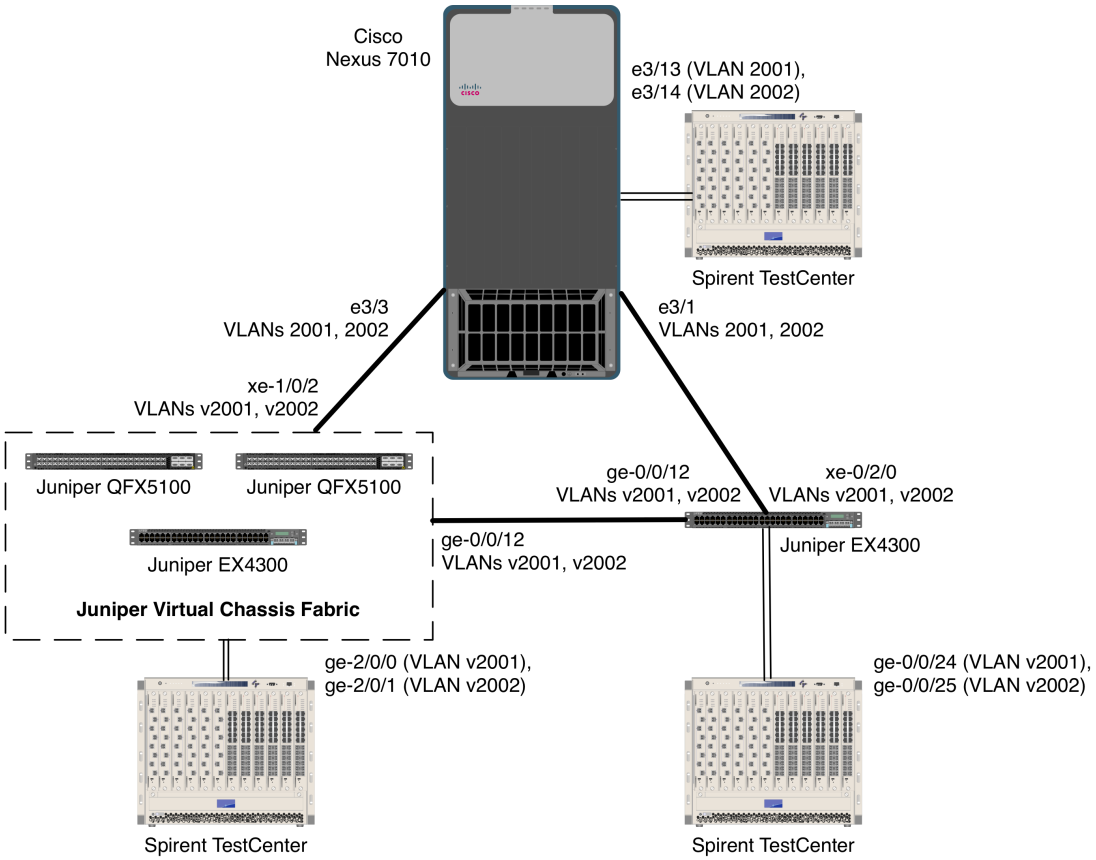

Figure 14: MSTP validation topology

Juniper commands

On both Juniper devices:

1. Create VLANs "v2001" and "v2002":

```
admin@VCF> config
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set vlans v2002 vlan-id 2002
```

2. Configure access-mode interfaces for VLANs v2001 and v2002, respectively. The Spirent TestCenter traffic generator/analyzer will attach to these ports.

On the Juniper Virtual Chassis Fabric:

```
admin@VCF# set interfaces ge-2/0/0 description "v2001 to Spirent"
admin@VCF# set interfaces ge-2/0/0.0 family ethernet-switching vlan members
v2001
admin@VCF# set interfaces ge-2/0/2 description "v2002 to Spirent"
admin@VCF# set interfaces ge-2/0/1.0 family ethernet-switching vlan members
v2002
```

On the Juniper EX4300:

```
admin@EX4300# set interfaces ge-0/0/24 description “v2001 to Spirent”
admin@EX4300# set interfaces ge-0/0/24.0 family ethernet-switching vlan
members v2001
admin@EX4300# set interfaces ge-0/0/25 description “v2002 to Spirent”
admin@EX4300# set interfaces ge-0/0/25.0 family ethernet-switching vlan
members v2002
```

3. Configure trunk ports that allow tagged traffic from VLANs v2001 and v2002.

On the Juniper Virtual Chassis Fabric:

```
admin@VCF# set interfaces xe-1/0/2 description “RSTP to Nexus 7010 e3/3”
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2002
admin@VCF# set interfaces ge-0/0/12 description “RSTP to Juniper EX4300 ge-
0/0/12”
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
v2002
```

On the Juniper EX4300:

```
admin@EX4300# set interfaces xe-0/2/0 description “RSTP to Nexus 7010 e3/1”
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members v2001
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members v2002
admin@EX4300# set interfaces ge-0/0/12 description “RSTP to Juniper EX4300
ge-0/0/12”
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members v2001
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members v2002
```

4. Enable multiple spanning tree. This requires deleting rapid spanning tree (if enabled) and then setting one multiple spanning tree instance (MSTI) per VLAN.

On the Juniper Virtual Chassis Fabric:

```
admin@VCF# delete protocols rstp
admin@VCF# set protocols mstp interface xe-1/0/2
admin@VCF# set protocols mstp interface ge-0/0/12
admin@VCF# set protocols mstp msti vlan v2001
```

```
admin@VCF# set protocols mstp msti vlan v2002
admin@VCF# commit
```

On the Juniper EX4300:

```
admin@VCF# delete protocols rstp
admin@VCF# set protocols mstp interface xe-0/2/0.0
admin@VCF# set protocols mstp interface ge-0/0/12.0
admin@VCF# set protocols mstp msti vlan v2001
admin@VCF# set protocols mstp msti vlan v2002
admin@VCF# commit
```

Cisco commands

1. Create VLANs 2001 and 2002:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001-2002
Nexus7010(config-vlan)# exit
```

2. Configure ports Ethernet3/13 and Ethernet3/14 as access-mode ports for VLANs 2001 and 2002, respectively:

```
Nexus7010(config)# interface Ethernet3/13
Nexus7010(config-if)# description MSTP to Spirent vlan 2001
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/14
Nexus7010(config-if)# description MSTP to Spirent vlan 2002
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2002
Nexus7010(config-if)# no shutdown
```

Note that Cisco Nexus 7000 switches require the explicit "**no shutdown**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

3. Configure ports Ethernet3/1 and Ethernet3/3 as trunk ports that allow tagged traffic for VLANs 2001 and 2002:

```
Nexus7010(config-if)# interface Ethernet3/1
Nexus7010(config-if)# description MSTP to Juniper VCF xe-1/0/2
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2002
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# interface Ethernet3/3
Nexus7010(config-if)# description MSTP to Juniper EX4300 xe-0/2/0
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2002
Nexus7010(config-if)# no shutdown
```

Note that Cisco Nexus 7000 switches require the explicit "**no shutdown**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

4. Enable multiple spanning tree. This requires defining multiple spanning tree as the mode of operation and adding one multiple spanning tree instance (MSTI) per VLAN.

```
Nexus7010(config)# spanning-tree mode mst
Nexus7010(config)# spanning-tree mst configuration
Nexus7010(config-mst)# instance 1 vlan 2001
Nexus7010(config-mst)# instance 2 vlan 2002
Nexus7010(config-mst)# end
```

Validation

The command **show spanning-tree bridge brief** will display a summary of spanning tree parameters:

```
admin@VCF> show spanning-tree bridge brief
STP bridge parameters
Context ID : 0
Enabled protocol : MSTI
...
```

The command **show spanning-tree interface brief** will display a summary of spanning tree parameters on a per-interface and per-MSTI basis.

```
admin@VCF> show spanning-tree interface brief

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
xe-1/0/2       128:491      128:387  32768.0024f7189ec1      2000  FWD    ROOT
ge-0/0/12      128:616      128:490  32768.100e7ea6ffc1     20000  BLK    ALT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
xe-1/0/40      128:493      128:493  32769.100e7eb06791      2000  FWD    DESG
ge-0/0/12      128:616      128:490  32769.100e7ea6ffc1     20000  BLK    ALT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
xe-1/0/2       128:491      128:387  32770.0024f7189ec1      2000  FWD    ROOT
ge-0/0/12      128:616      128:490  32770.100e7ea6ffc1     20000  BLK    ALT
```

The equivalent command for the Cisco Nexus 7010 is "**show spanning-tree <vlan ID>**". In this example, the Nexus 7010 is the root bridge for VLAN 2001.

```
Nexus7010# show spanning-tree vlan 2001


MST0001
  Spanning tree enabled protocol mstp
  Root ID    Priority    32769
             Address     0024.f718.9ec1
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     0024.f718.9ec1
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth3/1           Desg FWD 2000      128.385  P2p
Eth3/3           Desg FWD 2000      128.387  P2p
```

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source to each VLAN and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then rapid spanning tree convergence time is 47 ms.

Spanning tree case 3: VLAN spanning tree protocol (VSTP) and Per-VLAN Spanning Tree Plus (PVST+)

Objective

To verify interoperability of Juniper VLAN spanning tree protocol (VSTP) and Cisco per-VLAN spanning tree protocol plus (PVST+) between Juniper and Cisco switches, respectively.
To measure convergence time of a VSTP-PVST+ topology between Juniper and Cisco switches after link failure.

Background

With Juniper's VLAN Spanning Tree Protocol (VSTP), Juniper switches can run one or more spanning tree instances per VLAN. As stated in the Junos Software Guide, VSTP "enables more intelligent tree spanning, because each VLAN can have interfaces enabled or disabled depending on the paths available to that specific VLAN."

The goal of this exercise is to demonstrate interoperability in a multiple-VLAN environment using VSTP running on Juniper EX switches and PVST+ running on Cisco Nexus or Catalyst switches.

Topology

This example uses redundant links between two Juniper switches and one Cisco switch. VLAN IDs of 2001 and 2002 have been defined on both the Juniper and Cisco switches. VSTP is enabled on the Juniper switches, while the Cisco Nexus 7010 runs Rapid PVST+.

Figure 15 illustrates the VSTP-PVST+ validation test bed. This example involves a standalone Juniper EX4300 switch; a Juniper Virtual Chassis Fabric (comprised of two Juniper QFX5100 and one Juniper EX4300 switches); and a Cisco Nexus 7010.

As shown in the figure, the links interconnecting each switch are trunk ports that allow tagged traffic from VLAN IDs 2001 and 2002. Two access-mode ports are configured on each switch: One apiece for VLAN IDs 2001 and 2002. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology in each VLAN.

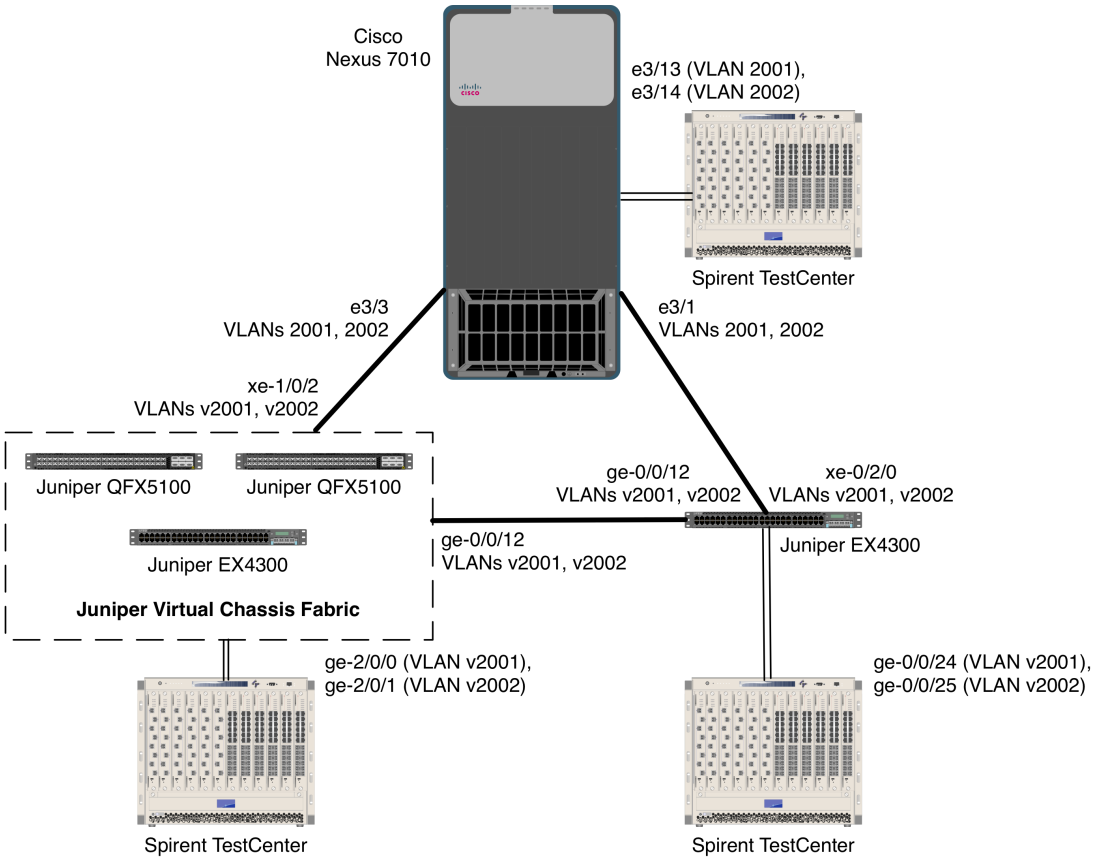

Figure 15: VSTP-PVST+ validation topology

Juniper commands

On both Juniper devices:

1. Create VLANs "v2001" and "v2002":

```
admin@VCF> config
admin@VCF# set vlans v2001 vlan-id 2001
admin@VCF# set vlans v2002 vlan-id 2002
```

2. Configure access-mode interfaces for VLANs v2001 and v2002, respectively. The Spirent TestCenter traffic generator/analyzer will attach to these ports.

On the Juniper Virtual Chassis Fabric:

```
admin@VCF# set interfaces ge-2/0/0 description "v2001 to Spirent"
admin@VCF# set interfaces ge-2/0/0.0 family ethernet-switching vlan members
v2001
admin@VCF# set interfaces ge-2/0/2 description "v2002 to Spirent"
admin@VCF# set interfaces ge-2/0/1.0 family ethernet-switching vlan members
v2002
```

On the Juniper EX4300:

```
admin@EX4300# set interfaces ge-0/0/24 description “v2001 to Spirent”
admin@EX4300# set interfaces ge-0/0/24.0 family ethernet-switching vlan
members v2001
admin@EX4300# set interfaces ge-0/0/25 description “v2002 to Spirent”
admin@EX4300# set interfaces ge-0/0/25.0 family ethernet-switching vlan
members v2002
```

3. Configure trunk ports that allow tagged traffic from VLANs v2001 and v2002.

On the Juniper Virtual Chassis Fabric:

```
admin@VCF# set interfaces xe-1/0/2 description “VSTP to Nexus 7010 e3/3”
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces xe-1/0/2 unit 0 family ethernet-switching vlan
members v2002
admin@VCF# set interfaces ge-0/0/12 description “VSTP to Juniper EX4300 ge-
0/0/12”
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
members v2001
admin@VCF# set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan
v2002
```

On the Juniper EX4300:

```
admin@EX4300# set interfaces xe-0/2/0 description “VSTP to Nexus 7010 e3/1”
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members v2001
admin@EX4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members v2002
admin@EX4300# set interfaces ge-0/0/12 description “VSTP to Juniper EX4300
ge-0/0/12”
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode trunk
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members v2001
admin@EX4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members v2002
```

4. Enable VLAN spanning tree protocol. This requires deleting rapid spanning tree or multiple spanning tree (if enabled) and then enabling VSTP for each VLAN.

On both Juniper devices:

```
admin@VCF# delete protocols rstp
admin@VCF# delete protocols mstp
admin@VCF# set protocols vstp vlan v2001
admin@VCF# set protocols vstp vlan v2002
admin@VCF# commit
```

Cisco commands

1. Create VLANs 2001 and 2002:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001-2002
Nexus7010(config-vlan)# exit
```

2. Configure ports Ethernet3/13 and Ethernet3/14 as access-mode ports for VLANs 2001 and 2002, respectively:

```
Nexus7010(config)# interface Ethernet3/13
Nexus7010(config-if)# description RPVST+ to Spirent vlan 2001
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/14
Nexus7010(config-if)# description RPVST+ to Spirent vlan 2002
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2002
Nexus7010(config-if)# no shutdown
```

Note that Cisco Nexus 7000 switches require the explicit "**`no shutdown`**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

3. Configure ports Ethernet3/1 and Ethernet3/3 as trunk ports that allow tagged traffic for VLANs 2001 and 2002:

```
Nexus7010(config-if)# interface Ethernet3/1
Nexus7010(config-if)# description RPVST+ to Juniper VCF xe-1/0/2
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2002
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# interface Ethernet3/3
Nexus7010(config-if)# description RPVST+ to Juniper EX4300 xe-0/2/0
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2002
Nexus7010(config-if)# no shutdown
```

Note that Cisco Nexus 7000 switches require the explicit "**`no shutdown`**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

2. Enable PVST Plus. On a new switch configuration, this step should be unnecessary since PVST Plus is enabled by default.

```
Nexus7010(config)# spanning-tree mode rstp
Nexus7010(config)# end
```

On Cisco Catalyst switches, the equivalent command is "**spanning-tree mode rapid-pvst**".

Validation

The command **show spanning-tree bridge brief** will display a summary of spanning tree parameters:

```
admin@EX4200> show spanning-tree bridge brief
STP bridge parameters
Context ID : 0
Enabled protocol : VSTP
...
```

The command **show spanning-tree interface brief** will display a summary of spanning tree parameters on a per-interface and per-VLAN basis:

```
admin@VCF> show spanning-tree interface brief

Spanning tree interface parameters for VLAN 2001

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
ge-0/0/12     128:4234      128:490  34769.100e7ea6ffc0     20000  BLK    ALT
xe-1/0/2      128:4236      128:387  34769.0024f7189ec1      2000  FWD    ROOT

Spanning tree interface parameters for VLAN 2002

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
ge-0/0/12     128:4234      128:490  34770.100e7ea6ffc0     20000  BLK    ALT
xe-1/0/2      128:4236      128:387  34770.0024f7189ec1      2000  FWD    ROOT
```

The equivalent command for the Cisco Nexus 7010 is "**show spanning-tree <vlan ID>**". In this example, the Nexus 7010 is the root bridge for VLAN 2001:

```
Nexus7010# show spanning-tree vlan 2001

LAN2001
  Spanning tree enabled protocol rstp
  Root ID    Priority    34769
             Address     0024.f718.9ec1
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    34769  (priority 32768 sys-id-ext 2001)
             Address     0024.f718.9ec1
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth3/1           Desg FWD 2         128.385  P2p
Eth3/3           Desg FWD 2         128.387  P2p
```

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source to each VLAN and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then rapid spanning tree convergence time is 47 ms.

Virtual LAN (VLAN) trunking

Objective

To verify interoperability of IEEE 802.1Q VLAN trunking between Juniper and Cisco switches using tagged traffic.
To verify interoperability of IEEE 802.1Q VLAN trunking between Juniper and Cisco switches using untagged (native) traffic.

Background

The IEEE 802.1Q specification defines a method for defining virtual broadcast domains. A 4-byte VLAN header, usually called a “tag,” allows definition of broadcast domains that may differ from physical switch topology. Without VLANs, all switch ports are members of the same broadcast domain; with VLAN tagging, a network manager can set up multiple broadcast domains across switches, and restrict broadcasts for different VLANs on different ports.

Topology

This configuration example will validate VLAN trunking interoperability between Juniper and Cisco switches in three ways:

- The switches will forward allowed tagged traffic from multiple VLANs across a trunk port
- The switches will forward allowed untagged traffic from a native VLAN across a trunk port
- The switches will not forward disallowed tagged traffic across a trunk port

The final example above is a negative test to verify that switches will forward only traffic explicitly permitted by VLAN trunking configurations.

Figure 16 below illustrates the test bed used to verify VLAN trunking operation. In this example, a VLAN trunk carries allowed VLAN traffic between a Juniper Virtual Chassis (comprising two Juniper EX9208 switches) and a Cisco Nexus 7010. Both switches use 10-gigabit Ethernet interfaces for the trunk port in this example, though VLAN trunking also would work on any matched pair of Ethernet interfaces. The trunk ports on each switch will allow tagged traffic with VLAN IDs 2001 and 2002, and untagged ("native") traffic from ports with a VLAN ID of 2003. A fourth VLAN, with a ID of 2004, also exists, but the trunk port is configured not to allow that traffic.

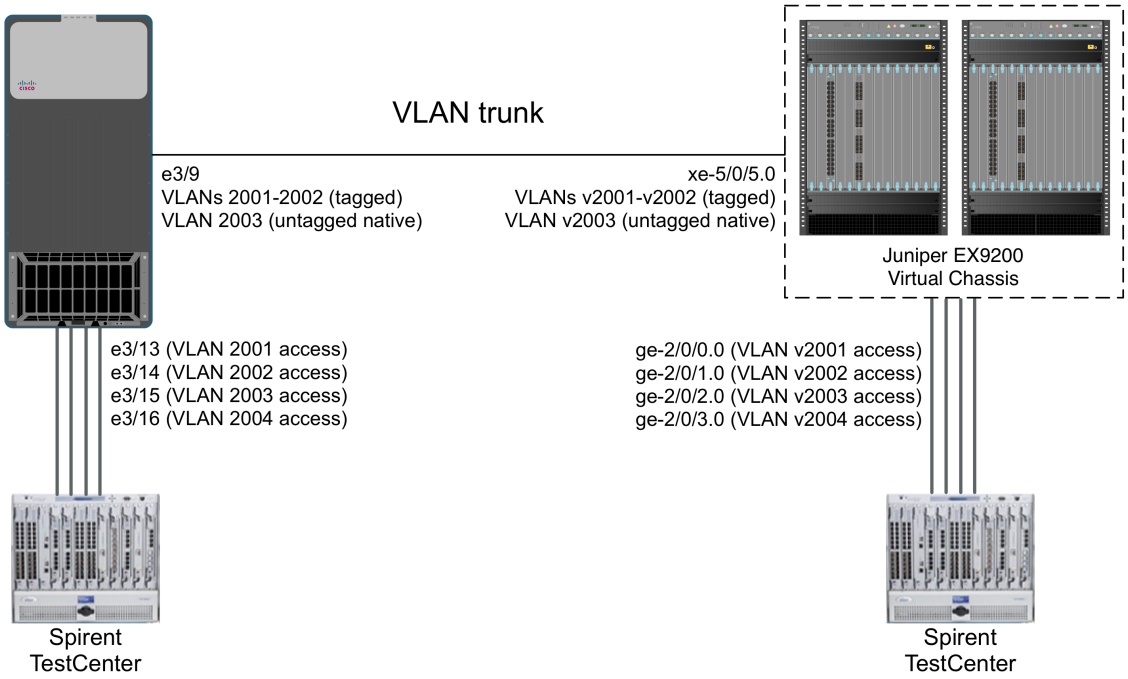

Figure 16: VLAN trunking validation topology

Juniper commands

1. Define VLANs v2001 through v2004 with VLAN IDs of 2001 through 2004 respectively:

```
admin@EX9208> config
admin@EX9208# set vlans v2001 vlan-id 2001
admin@EX9208# set vlans v2002 vlan-id 2002
admin@EX9208# set vlans v2003 vlan-id 2003
admin@EX9208# set vlans v2004 vlan-id 2004
admin@EX9208# set vlans v2005 vlan-id 2005
```

2. Define a VLAN trunk port that allows tagged traffic from VLANs v2001 and v2002 and untagged traffic from VLAN v2003:

```
admin@EX9208# set interfaces xe-5/0/5 description "VLAN trunk to Nexus 7010
e3/9"
admin@EX9208# set interfaces xe-5/0/5 unit 0 family ethernet-switching vlan
members v2001
admin@EX9208# set interfaces xe-5/0/5 unit 0 family ethernet-switching
interface-mode trunk
admin@EX9208# set interfaces xe-5/0/5 unit 0 family ethernet-switching vlan
members v2001
admin@EX9208# set interfaces xe-5/0/5 unit 0 family ethernet-switching vlan
members v2002
admin@EX9208# set interfaces xe-5/0/5 unit 0 family ethernet-switching vlan
members v2003
```

3. On the VLAN trunk port, allow native untagged traffic from VLAN v2003. Note that the **`native-vlan-id`** command takes the VLAN ID and not the VLAN name as an argument:

```
admin@EX9208# set interfaces xe-5/0/5 native-vlan-id 2003
```

4. Define access-mode interfaces allowing untagged traffic from VLANs v2001 through v2004:

```
admin@EX9208# set interfaces ge-2/0/0 description "to stc v2001"
admin@EX9208# set interfaces ge-2/0/0 unit 0 family ethernet-switching
interface-mode access
admin@EX9208# set interfaces ge-2/0/0 unit 0 family ethernet-switching vlan
members v2001
admin@EX9208# set interfaces ge-2/0/1 description "to stc v2002"
admin@EX9208# set interfaces ge-2/0/1 unit 0 family ethernet-switching
interface-mode access
admin@EX9208# set interfaces ge-2/0/1 unit 0 family ethernet-switching vlan
members v2002
admin@EX9208# set interfaces ge-2/0/2 description "to stc v2003"
admin@EX9208# set interfaces ge-2/0/2 unit 0 family ethernet-switching
interface-mode access
admin@EX9208# set interfaces ge-2/0/2 unit 0 family ethernet-switching vlan
members v2003
admin@EX9208# set interfaces ge-2/0/3 description "to stc v2004"
admin@EX9208# set interfaces ge-2/0/3 unit 0 family ethernet-switching
interface-mode access
```

```
admin@EX9208# set interfaces ge-2/0/3 unit 0 family ethernet-switching vlan
members v2004
admin@EX9208# commit
```

Cisco commands

1. Define VLANs 2001 through 2004:

```
Nexus7010# configure terminal
Nexus7010(config)# vlan 2001-2004
Nexus7010(config-vlan)# exit
```

2. Define a VLAN trunk port that allows tagged traffic from VLANs 2001 and 2002 and native untagged traffic from VLAN 2003:

```
Nexus7010(config)# interface Ethernet3/9
Nexus7010(config-if)# description VLAN trunk to EX9208 xe-5/0/5
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport mode trunk
Nexus7010(config-if)# switchport trunk native vlan 2003
Nexus7010(config-if)# switchport trunk allowed vlan 2001-2003
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# exit
```

Note that Cisco Nexus 7000 switches require the explicit "**`no shutdown`**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

3. Define access-mode interfaces allowing untagged traffic from VLANs 2001 through 2003:

```
Nexus7010(config)# interface Ethernet3/13
Nexus7010(config-if)# description to stc vlan 2001
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2001
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/14
Nexus7010(config-if)# description to stc vlan 2002
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2002
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/15
Nexus7010(config-if)# description to stc vlan 2003
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2003
Nexus7010(config-if)# no shutdown
Nexus7010(config)# interface Ethernet3/16
Nexus7010(config-if)# description to stc vlan 2004
Nexus7010(config-if)# switchport
Nexus7010(config-if)# switchport access vlan 2004
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Cisco Catalyst switches also will require the "**`switchport mode access`**" command on each access-mode interface.

Note that Cisco Nexus 7000 switches require the explicit "**`no shutdown`**" command to enable interfaces. This command is not required in Cisco Catalyst switches.

Validation

The Spirent TestCenter traffic generator/analyzer can be configured to offer bidirectional traffic between pairs of access-mode interfaces on each switch. In all cases – involving unicast, broadcast or multicast traffic – traffic will stay local to the VLAN in which it is defined. For example, traffic offered to VLAN v2001 on the Juniper switch will be forwarded only to interfaces in VLAN 2001 on the Cisco switch and vice-versa.

If desired, port mirroring can be enabled on either switch to verify that the trunk ports carry tagged traffic for VLAN IDs 2001 and 2002 and untagged traffic for VLAN ID 2003.

As a final verification that VLANs limit broadcast domains, Spirent TestCenter can be configured to offer traffic to the access ports with an VLAN ID of 2004. The trunk ports on both switches will not forward this traffic.

Virtual Router Redundancy Protocol (VRRP) interoperability

Objective

To validate failover functionality of the virtual router redundancy protocol (VRRP) between Juniper and Cisco switches configured as routers.

Background

As described in RFC 5798, two or more routers can make use of VRRP to enhance network availability. With VRRP, multiple routers share a single virtual IP address. One router acts as the master (active) device, while all others act as backups. If the master router fails (or if a link fails on the interface configured with the virtual IP address), one of the backup routers takes over as master.

Topology

In this example, Juniper and Cisco switches are both configured to route IP traffic. The interfaces connecting the switches each have unique IP addresses configured – 192.18.38.1/24 for the Juniper Virtual Chassis and 192.18.38.2/24 for the Cisco Nexus 7010. The Juniper and Cisco devices also share a single virtual IP address of 192.18.38.254/24, with the Juniper device initially acting as VRRP master.

The PCs attached to the Juniper and Cisco devices each use the virtual IP address of 192.18.38.254/24 as their default gateway. In the event of a failure of the master (Juniper) device, the backup (Cisco) device will take over as master.

Figure 17 illustrates the VRRP validation test bed. In this example, both the Juniper and Cisco devices advertise the virtual IP address of 192.18.38.254. On both the Juniper and Cisco devices, the IP address is assigned to the physical interface rather than a VLAN interface. However, VRRP would work equally well with IP addresses assigned to VLAN interfaces.

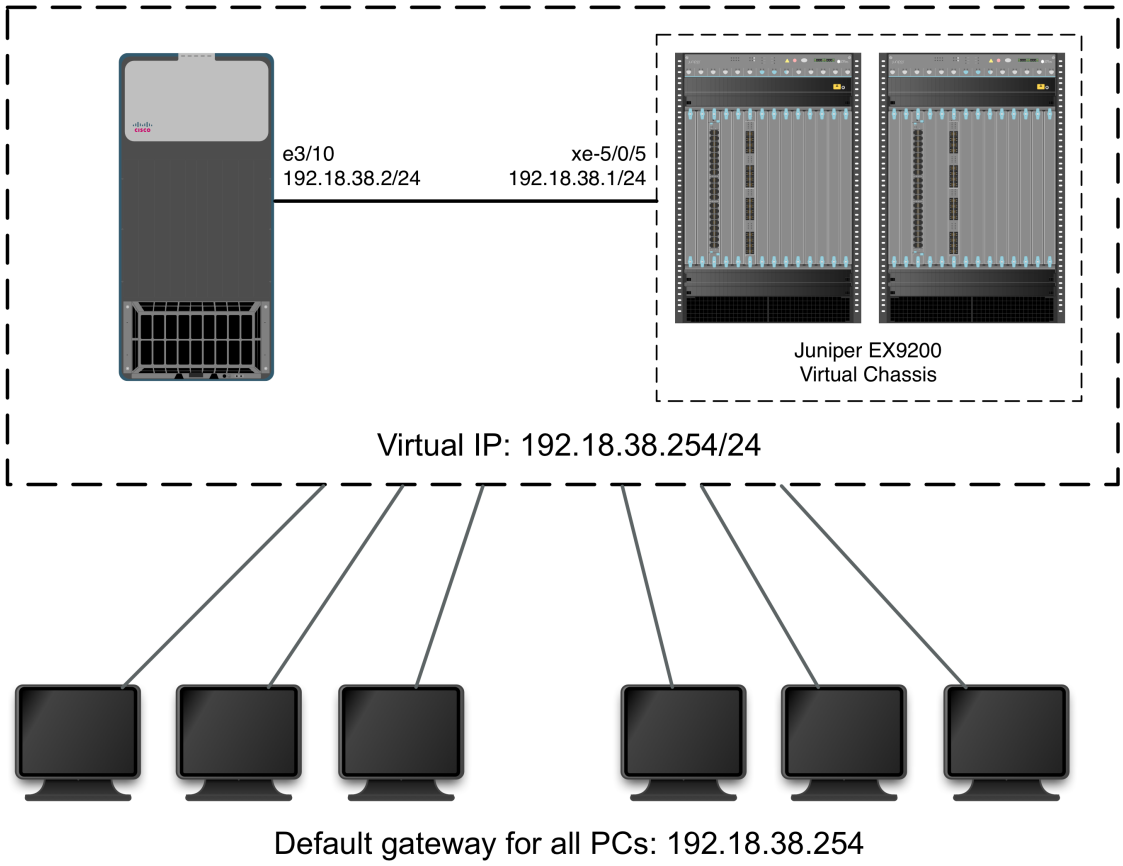

Figure 17: VRRP validation topology

Juniper commands

Configure an IP address on interface xe-5/0/5, and also define **`vrrp-group 1`** on that interface:

```
admin@EX9208> config
admin@EX9208# set interfaces xe-5/0/5 description "VRRP link to Nexus 7010
e3/10"
admin@EX9208# set interfaces xe-5/0/5 unit 0 family inet address 192.18.38.1/24
vrrp-group 1 virtual-address 192.18.38.254
admin@EX9208# set interfaces xe-5/0/5 unit 0 family inet address 192.18.38.1/24
vrrp-group 1 priority 254
admin@EX9208# set interfaces xe-5/0/5 unit 0 family inet address 192.18.38.1/24
vrrp-group 1 preempt
```

```
admin@EX9208# set interfaces xe-5/0/5 unit 0 family inet address 192.18.38.1/24
vrrp-group 1 accept-data
admin@EX9208# commit
```

The optional **priority 255** statement forces the Juniper switch's virtual IP address to become the master VRRP instance, assuming the Cisco Nexus 7010 device uses a lower priority value. The legal range of VRRP priorities is 1 though 255, with 255 being highest.

Cisco commands

1. Ensure the VRRP feature is installed. This step is not required with Cisco Catalyst switches:

```
Nexus7010# configure terminal
Nexus7010(config)# feature vrrp
```

2. Define interface and VRRP virtual IP addresses on interface Ethernet3/10:

```
Nexus7010(config)# interface Ethernet3/10
Nexus7010(config-if)# description VRRP to EX9208 xe-5/0/5
Nexus7010(config-if)# no switchport
Nexus7010(config-if)# ip address 192.18.38.2/24
Nexus7010(config-if)# vrrp 1
Nexus7010(config-if-vrrp)# address 192.18.38.254
Nexus7010(config-if-vrrp)# no shutdown
Nexus7010(config-if-vrrp)# exit
Nexus7010(config-if)# no shutdown
Nexus7010(config-if)# end
```

Note the explicit "**no shutdown**" commands for both physical and virtual interfaces. These commands are mandatory with Cisco Nexus 7000 switches, but are not required with Cisco Catalyst switches.

Validation

On Juniper devices, the **show vrrp summary** command will indicate the current VRRP state on each system. In the following examples, the Juniper Virtual Chassis acts as VRRP master and the Cisco Nexus 7010 acts as a backup.

On the Juniper Virtual Chassis:

```
admin@EX9208# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
xe-5/0/5.0    up              1   master          Active     lcl
192.18.38.1
                                                             vip

192.18.38.254
```

The equivalent command on Cisco devices is **show vrrp summary**:

```
Nexus7010# show vrrp
      Interface  VR IpVersion Pri   Time Pre State   VR IP addr
------------------------------------------------------------------
   Ethernet3/10   1   IPV4    100     1 s  Y  Backup 192.18.38.254
```

Note that both devices agree the master router is 192.18.38.1 (on the Juniper Virtual Chassis), and both use a virtual IP address of 192.18.38.254.

If a router or link fails, the backup router should take over as the master. In this example, the Cisco Nexus 7010 is promoted to master statement by reducing the Juniper Virtual Chassis' priority to 10. Since the Cisco Nexus 7010 uses a priority of 250 by default, it will take over as master once this Juniper Virtual Chassis configuration change is committed:

```
admin@EX9208> configure
admin@EX9208# set interfaces xe-5/0/5 family inet address 192.18.38.1/24
vrrp-group 1 virtual-address 192.18.38.254 priority 10
admin@EX9208# commit
```

After this change, the Juniper device becomes the backup router, and the master router is now 192.18.38.254 on the Cisco device:

```
admin@EX9208# run show vrrp summary
Interface     State       Group   VR state       VR Mode    Type   Address
xe-5/0/5.0    up              1   backup          Active    lcl
192.18.38.1
                                                            vip
192.18.38.254
```

The Cisco device agrees that it is now the VRRP master:

```
Nexus7010# show vrrp summary
VRRP Summary
------------
Total Number of Groups Configured: 1
        Init : 0        Backup : 0      Master : 1
```

Wi-Fi passthrough

Objective

To verify the ability of a Juniper switch to forward Wi-Fi management traffic between a Cisco Wi-Fi controller and Cisco Wi-Fi access points.

Background

Many enterprises provision wireless LAN access using one or more centrally managed controllers that monitor and manage RF and data networking parameters in real time. In this model, Wi-Fi access points (APs) distributed throughout the enterprise depend on the controllers for their configurations. Thus, reliable connectivity between controllers and access points is essential.

Juniper switches can carry Cisco and other vendors' Wi-Fi traffic transparently, with no special configuration required.

Topology

As shown in Figure 18, a centrally located Cisco 5508 Wi-Fi controller in a data center must communicate with Cisco 3602 and Cisco 3702 APs in a campus network. The Cisco controller connects with a Juniper EX9208 core switch. The Cisco APs connect with a Juniper EX4300 access switch using the Power Over Ethernet Plus (PoE+) specification.

The Cisco controller and APs both use the **`default`** VLAN, which here uses a VLAN ID of 1. Trunk ports between the Juniper switches allow traffic from all VLANs.

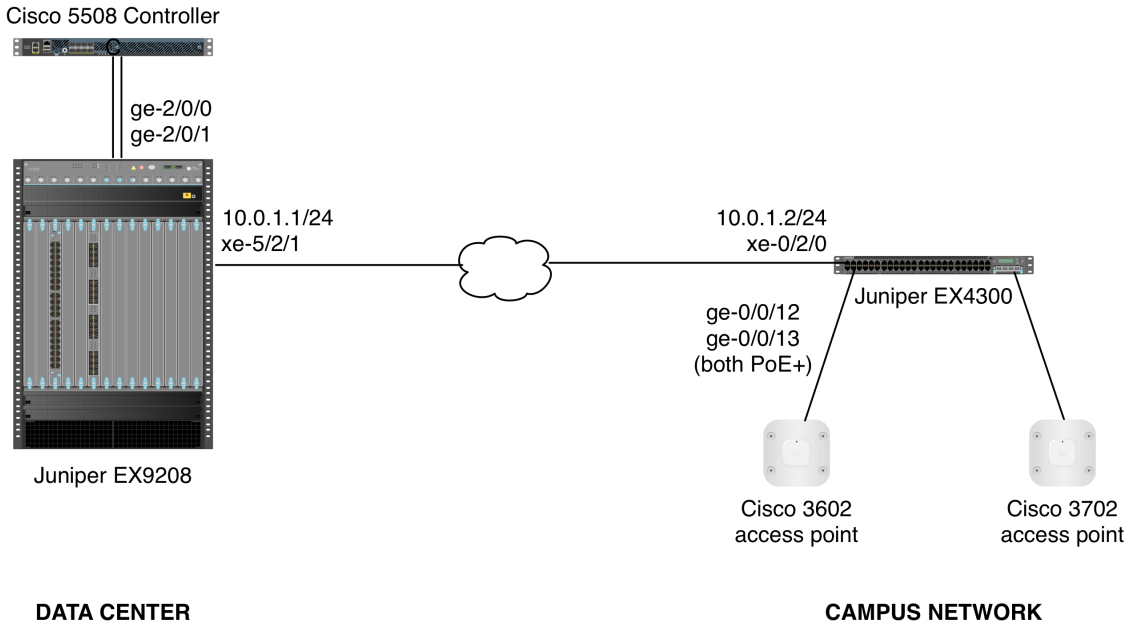

Figure 18: Wi-Fi passthrough validation topology

Juniper commands

No special configuration is needed on the Juniper switches. In this example, both access and trunk ports are members of the default VLAN, and the trunk port between switches allows traffic from all VLANs.

On the Juniper EX9208 switch:

1. Create a VLAN called **`default`** and assign a VLAN ID of 1:

```
admin@ex9208> configure
admin@ex9208# set vlans default vlan-id 1
```

2. (Optional) Create a Layer-3 interface for the **default** VLAN and assign an IP address to that interface. This step is not mandatory for Wi-Fi passthrough operation:

```
admin@ex9208# set vlans default l3-interface irb.1
admin@ex9208# set interfaces irb unit 1 family inet address 10.0.1.1/24
```

3. Configure the inter-switch link as a VLAN trunk and allow traffic from all VLANs:

```
admin@ex9208# set interfaces xe-5/2/1 description "trunk to EX4300 xe-
0/2/0"
admin@ex9208# set interfaces xe-5/2/1 unit 0 family ethernet-switching
interface-mode trunk
admin@ex9208# set interfaces xe-5/2/1 unit 0 family ethernet-switching vlan
members all
```

This example uses the Junos keyword **all** to permit all VLAN traffic across the trunk. If desired, forwarding can be restricted to specific VLANs. For example, this command would permit traffic only from the **default** VLAN:

```
admin@ex9208# set interfaces xe-5/2/1 unit 0 family ethernet-switching vlan
members default
```

4. Configure two interfaces to allow traffic from the Cisco controller, both as access-mode members of the **default** VLAN:

```
admin@ex9208# set interfaces ge-2/0/0 description "to Cisco 5508 Wi-Fi
controller"
admin@ex9208# set interfaces ge-2/0/0 unit 0 family ethernet-switching
interface-mode access
admin@ex9208# set interfaces ge-2/0/0 unit 0 family ethernet-switching vlan
members default
admin@ex9208# set interfaces ge-2/0/1 description "to Cisco 5508 Wi-Fi
controller"
admin@ex9208# set interfaces ge-2/0/1 unit 0 family ethernet-switching
interface-mode access
admin@ex9208# set interfaces ge-2/0/1 unit 0 family ethernet-switching vlan
members default
admin@ex9208# commit
```

On the Juniper EX4300 switch:

1. Create a VLAN called **default** and assign a VLAN ID of 1:

```
admin@ex4300> configure
admin@ex4300# set vlans default vlan-id 1
```

2. (Optional) Create a Layer-3 interface for the **default** VLAN and assign an IP address to that interface. This step is not mandatory for Wi-Fi passthrough operation:

```
admin@ex4300# set vlans default l3-interface irb.1
admin@ex4300# set interfaces irb unit 1 family inet address 10.0.1.2/24
```

3. Configure the inter-switch link as a VLAN trunk and allow traffic from all VLANs:

```
admin@ex4300# set interfaces xe-0/2/0 description "trunk to EX9208 xe-
5/2/1"
admin@ex4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching
interface-mode trunk
admin@ex4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members all
```

This example uses the Junos keyword **all** to permit all VLAN traffic across the trunk. If desired, forwarding can be restricted to specific VLANs. For example, this command would permit traffic only from the **default** VLAN:

```
admin@ex4300# set interfaces xe-0/2/0 unit 0 family ethernet-switching vlan
members default
```

4. Configure two interfaces to allow traffic from the Cisco APs, both as access-mode members of the **default** VLAN:

```
admin@ex4300# set interfaces ge-0/0/12 description "to Cisco AP 3602"
admin@ex4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
interface-mode access
admin@ex4300# set interfaces ge-0/0/12 unit 0 family ethernet-switching
vlan members default
admin@ex4300# set interfaces ge-0/0/13 description "to Cisco AP3702"
admin@ex4300# set interfaces ge-0/0/13 unit 0 family ethernet-switching
interface-mode access
admin@ex4300# set interfaces ge-0/0/13 unit 0 family ethernet-switching
vlan members default
```

5. Enable PoE+ on the interfaces to which the Cisco APs are attached:

```
admin@ex4300# set poe interface ge-0/0/12
admin@ex4300# set poe interface ge-0/0/13
admin@ex4300# commit
```

If desired, PoE+ can be enabled on all interfaces using Junos' **all** keyword:

```
admin@ex4300# set poe interface all
admin@ex4300# commit
```

Cisco commands

Configuration of the Cisco Wi-Fi controllers involves numerous RF as well as data networking parameters, and is beyond the scope of this document. For more information, consult the Cisco Wireless LAN Configuration Guide.

Validation

On the Cisco Wi-Fi controller attached to the Juniper EX9208, the command "**show ap summary**" will display a list of all APs managed by the controller.

```
(Cisco Controller) >show ap summary

Number of APs.................................... 2

Global AP User Name.............................. Not Configured
Global AP Dot1x User Name........................ Not Configured

AP Name             Slots  AP Model              Ethernet MAC
Location          Country  IP Address       Clients
------------------  -----  -------------------  -----------------  -------
---------  -------  --------------  -------
AP1                   2     AIR-CAP3602I-A-K9     c4:64:13:c0:7c:da  default
location  US       10.0.1.53        3
APb838.61a6.8530      2     AIR-CAP3702I-A-K9     b8:38:61:a6:85:30  default
location  US       10.0.1.59        5
```

Again on the Cisco controller, the command "**show client summary**" will display the total number of clients associated with the APs.

```
(Cisco Controller) show client summary

Number of Clients................................ 7
Number of PMIPV6 Clients......................... 0

                                                                  GLAN/
                                                                  RLAN/
MAC Address       AP Name           Slot Status        WLAN  Auth Protocol
Port Wired PMIPV6 Role
----------------- ----------------- ---- ------------ ----- ---- ---------
------- ---- ----- ------ ----------------
24:77:03:80:39:4c APb838.61a6.8530   0   Associated     1    No
802.11n(2.4 GHz) 1    No    No     Unassociated
24:77:03:80:4f:50 APb838.61a6.8530   0   Associated     1    No
802.11n(2.4 GHz) 1    No    No     Unassociated
7c:d1:c3:8b:96:a8 AP1                0   Associated     1    Yes
802.11n(2.4 GHz) 1    No    No     Local
7c:d1:c3:8b:b0:66 AP1                0   Associated     1    Yes
802.11n(2.4 GHz) 1    No    No     Local
7c:d1:c3:8c:ef:92 APb838.61a6.8530   0   Associated     1    Yes
802.11n(2.4 GHz) 1    No    No     Local
7c:d1:c3:8d:95:20 APb838.61a6.8530   0   Associated     1    Yes
802.11n(2.4 GHz) 1    No    No     Local
7c:d1:c3:8d:de:c8 AP1                0   Associated     1    Yes
802.11n(2.4 GHz) 1    No    No     Local
```

The command-line interface of the APs can verify correct PoE+ operation, including the ability of the Juniper EX4300 switch to supply power greater than the 15.4-watt limit of standard PoE.

```
APb838.61a6.8530>show controllers dot11Radio 0 powerreg

### Tx Power: Dot11 Config
        Serving Channel           11 (Type 0)
        Active Level Index        1 (Unitless)
        Active Level              23 dBm (OFDM 23 dBm)
        IEEE MIB format           TRUE
```

```
        Metric unit               dBm
        HD Active                 127
        Low Power Mode            Set High (Active: Set High)
        Cookie max                17
        Client Power              Max 23, Active Max 23 (Symmetric)
        Regulatory Limit          30
...
```

Note that the PoE+-capable AP draws 23 watts of power, well above the 15.4-watt limit of standard PoE.

Appendix A: Sample Configuration Files

This appendix lists URLs for the Juniper and Cisco switch configuration files used to verify interoperability. These files are freely available for download from a public Network Test server.

A copy of this document, a brief interoperability report and all Juniper and Cisco configuration files are available at http://networktest.com/jnpriop14.

Appendix B: Software Versions Tested

This appendix lists the software versions used for all test bed devices.

Juniper EX9208 (in Virtual Chassis and standalone configurations): Junos 13.3R1.6
Juniper QFX5100 (in Virtual Chassis Fabric): Junos 13.2-20140409_x_132_x51
Juniper EX4300 (in standalone configuration): Junos 13.2X51-D15.5
Juniper MX80: Junos 13.3R1.6
Cisco Catalyst 3850: IOS-XE 03.02.01.SE
Cisco Nexus 7010: NX-OS 6.1(4a)
Cisco Catalyst 6509: IOS 12.2(33)SXH1
Cisco 5500 Wi-Fi controller: 7.6.110.0
Spirent TestCenter: 4.33.0086

Appendix C: Disclaimer

Network Test Inc. has made every attempt to ensure that all test procedures were conducted with the utmost precision and accuracy, but acknowledges that errors do occur. Network Test Inc. shall not be held liable for damages which may result for the use of information contained in this document.