Using Layer 3 and vPC on the Cisco Nexus 5500 Series Device

This chapter describes virtual port channel (vPC) operations when Layer 3 routing features are enabled on the Cisco Nexus 5500 Series device.

This chapter includes the following sections:

vPC and First Hop Redundancy Protocol

When you use a Cisco Nexus 5548 switch or Cisco Nexus 5596UP switch as a default gateway for hosts, you can deploy the First Hop Redundancy Protocol (FHRP) to provide default gateway redundancy. Beginning with Cisco NX-OS Release 5.0(3)N1(1b), an active FHRP peer and a standby peer can perform Layer 3 forwarding when you enable vPC. This optimization improves bandwidth, avoids sending the Layer 3 traffic over the vPC peer link, and requires no configuration or protocol change. Only the FHRP active peer answers ARP requests. Because both active and standby FHRP peers can forward Layer 3 traffic, you do not need to configure an aggressive timer for FHRP to provide faster failover and convergence time if an active FHRP peer fails.

Figure 1-1 shows that the Layer 3 traffic that originated from the host and is destined to a host several hops away can be routed by both the Host Standby Router Protocol (HSRP) active and the HSRP standby switch.

Figure 1-1 vPC and FHRP

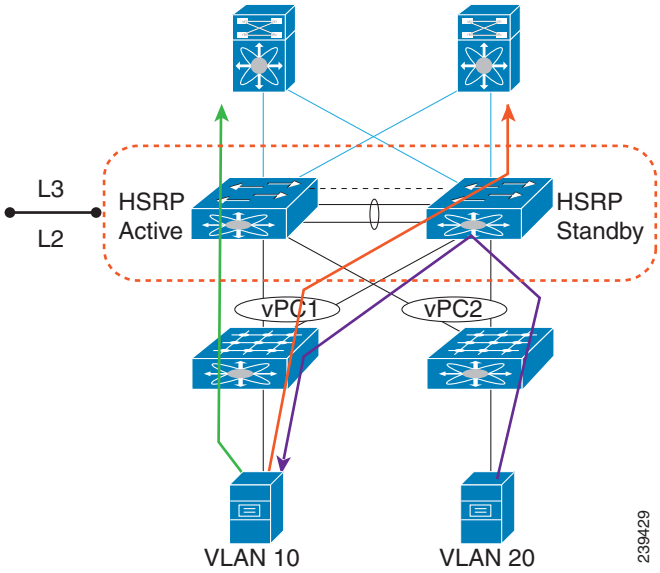

ARP Processing with vPC

When the host connects to a Cisco Nexus 5500 Platform switch and Cisco Nexus 2000 Fabric Extenders in a vPC topology, the host can send an ARP request to the FHRP standby peer due to a hashing algorithm. The ARP request that is received by the standby peer is forwarded to the active peer and the active peer can answer it with an ARP reply.

Similarly, when traffic is moving from north to south, such as when one Cisco Nexus 5500 Platform switch sends an ARP request to a host, the ARP reply might be sent to another switch. In such a case, the ARP reply is forwarded as a Layer 2 frame to the Cisco Nexus 5500 Platform switch that originated the ARP request.

As of Cisco NX-OS Release 5.0(3)N1(1b), ARP synchronization does not occur between two Cisco Nexus 5500 Platform switches. The two switches resolve and maintain their ARP table independently. When one vPC peer switch is reloaded, the switch needs to resolve the ARP by sending ARP requests to the hosts.

Layer 3 Forwarding for Packets to a Peer Switch MAC Address

Typically, a router performs a Layer 3 route table lookup and Layer 3 forwarding when the destination MAC in the Ethernet frame matches its own MAC address. Otherwise, the packets are switched (if Layer 2 functionality is enabled) or dropped. In a topology with Layer 3 and vPC enabled, a vPC peer switch could receive IP packets with the peer's MAC address as the destination MAC rather than the virtual MAC address (when FHRP is enabled) or its own MAC address. In this scenario, a Cisco Nexus 5500 Platform switch can forward the traffic to the peer using a peer link and the peer switch performs the Layer 3 forwarding.

The above scenario often happens with some filers. In the case of filers, they may achieve improved load balance and better performance by forwarding traffic to the Burnt-in-Address (BIA) of the routers instead of the HSRP MAC.

Figure 1-2 shows that when the NAS filer sends out packets with N5k-1's MAC RMAC-A as the destination MAC, the packets can be sent over to the N5k-2 switch due to the port channel hashing.

Figure 1-2 vPC and Peer-Gateway

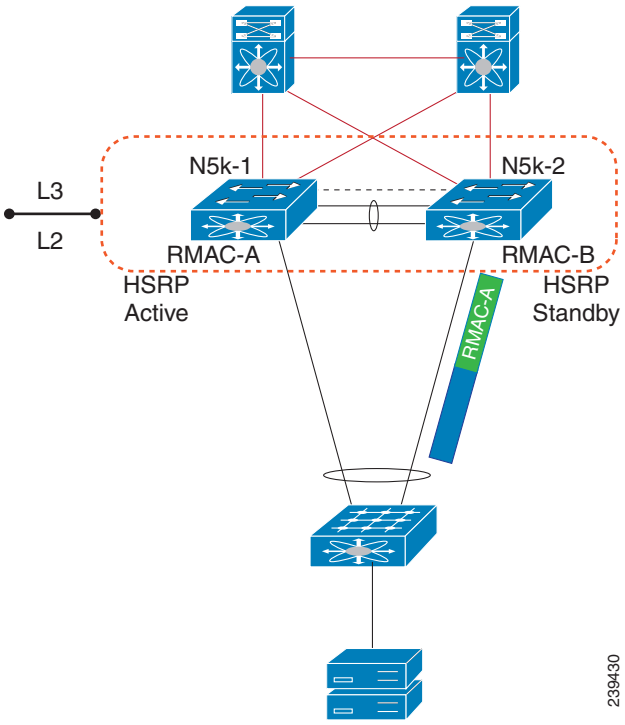

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), you can use the **peer-gateway** command to allow Cisco Nexus 5500 Platform switches to perform Layer 3 forwarding if the destination MAC of the incoming packet is the MAC of its vPC peer switch. The **peer-gateway** command avoids forwarding such packets to the vPC peer link.

Note You must configure the **peer-gateway** command on both vPC peer switches.

Improved Convergence with a vPC Topology and Layer 3 Routing

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), a delay timer was introduced to avoid the situation where a vPC member port is brought up before the Layer 3 is converged. For example, when one Cisco Nexus 5500 Platform switch is reloaded, the switch starts to receive traffic from hosts once the vPC member ports are up. A delay might occur before the switch establishes a routing protocol adjacency and learns all routes. During this period of the time, received traffic is dropped due to the lack of a route-to-destination address. Figure 1-3 shows an example of where the delay can be used to avoid black hole traffic when a Cisco Nexus 5000 Platform switch is configured for Layer 3 with vPC.

Figure 1-3 vPC Delay Restore

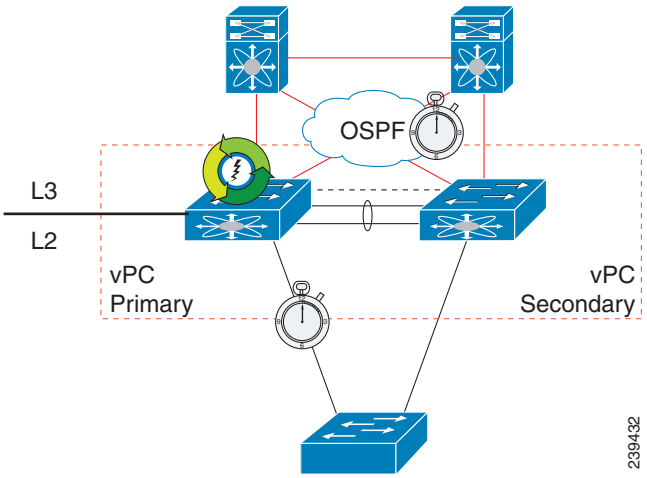

The delay restore feature allows you to configure a timed delay before vPC member ports are brought online. The delay allows the switch to learn all routes, to bring up the vPC member ports, and to forward traffic from hosts. The following example shows how to configure a timed delay of 120 seconds:

```
layer3-switch(config-vpc-domain)# delay restore ?
  <1-3600>  Delay in bringing up the vPC links (in seconds)
layer3-switch(config-vpc-domain)# delay restore 120
layer3-switch(config-vpc-domain)#
```

vPC Peer Link Failure

In addition to suspending vPC member ports, the vPC secondary switch also suspends its switched virtual interface (SVIs) when a vPC peer link is lost. When this occurs, the vPC secondary switch stops advertising the local subnets, which prevents traffic blackholing.

Layer 3 Module Failure

When a Layer 3 module fails on a Cisco Nexus 5500 Platform switch all Layer 3 interfaces are suspended, including Layer 3 port channel and SVI interfaces. As a result, the Layer 3 routing table on the neighboring routers is updated which results in the north to south traffic to be directed towards the peer Nexus 5500 Platform switch. The Layer 2 interfaces, including the Layer 2 port channel and out-of-band management interfaces, remain up.

In a non-vPC topology, when the Layer 3 and SVI interfaces are down, the redundant Cisco Nexus 5500 Platform switch becomes the active peer for all FHRP groups and it continues to forward traffic.

In a vPC topology, although the SVI interfaces are suspended, the vPC member ports are still up on the Cisco Nexus 5500 Platform switch. Even if the switch has a faulty Layer 3 module, Layer 2 traffic forwarding continues.

Figure 1-4 shows a topology where the Layer 3 module on N5k-2 fails. In this scenario, the Layer 3 connection toward the Layer 3 network and all SVI interfaces are suspended. However, the traffic from the hosts can still be sent to N5k-2 depending on the hash results. With the failure of the Layer 3 module, N5k-2 functions as a Layer 2 switch. It forwards the traffic to N5k-1, which forwards the traffic to the Layer 3 network. The return traffic is sent to N5k-1, which sends the traffic directly to the hosts.

Figure 1-4 Layer 3 Module Failure

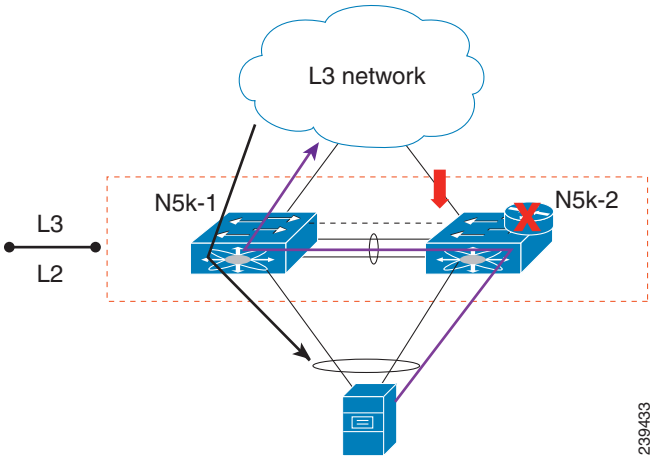

Note Only the Layer 3 traffic needs to cross the peer link. The VLAN traffic is switched by N5k-2 locally.

The peer gateway is disabled on both vPC switches if the Layer 3 module fails on one switch.

For topologies with in-band management, the failure of a Layer 3 module means that the connectivity to the management network and the management system is also lost.

Connecting to a Router in a vPC Topology

When you connect a router to a pair of Cisco Nexus 5500 Platform switches in a vPC topology and enable routing, traffic forwarding may result in suboptimal traffic paths crossing the peer link similar to the situation described in the "Layer 3 Forwarding for Packets to a Peer Switch MAC Address" section on page 1-2. We recommend that you use Layer 3 links for connections between the router and the Nexus 5500 switch, instead of a port channel with an IP address.

Note You cannot have a link for non-vPC traffic in parallel with a vPC topology. This can cause errors with the traffic forwarding logic resulting in duplicate or missed packets.

Figure 1-5 illustrates the topology that is not recommended. In this topology, control protocol packets may be hashed by the port channel to the wrong Cisco Nexus 5500 Platform switch, which would then forward the control packets to the correct routing peer (1.1.1.1) in the picture.

Figure 1-5 Control Traffic Forwarding in a vPC Topology

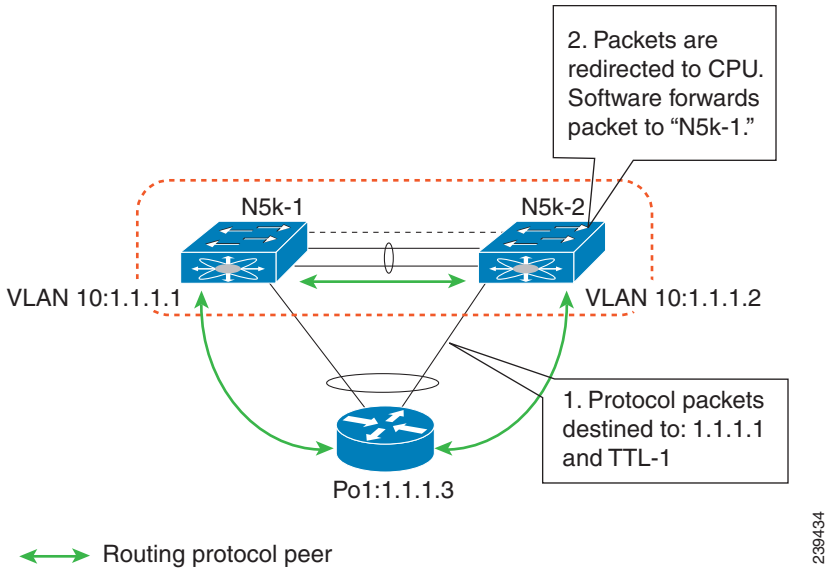

In this topology, we recommend that you use Layer 3 interfaces instead of vPC interfaces to connect routers to Cisco Nexus 5500 Platform switches whenever possible.

Figure 1-6, shows the recommended topology for connectivity of routers to a vPC domain. The router connects with Layer 3 interfaces 1.1.1.2 and 2.2.2.2 to the two vPC peers and these interfaces are not part of a vPC port channel.

Figure 1-6 Connecting a Router to a vPC Domain Using Layer 3 Interfaces

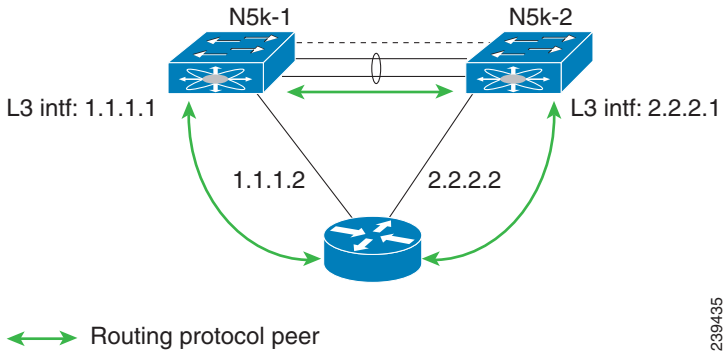

Dedicated VRF For a Keepalive Interface

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform switch supports VRF lite with a Layer 3 module and Enterprise license and you can create a VRF and assign the interface to a VRF. Prior to this release, two VRFs were created by default: the VRF management and VRF default. The management interface(mgmt0) and all SVI interfaces resided in the VRF management and VRF default respectively.

We recommend that you use an out-of-band management interface (mgmt0) as a vPC keepalive interface although you have the option to use the front-panel data port as a vPC keepalive interface. When you choose to use the front panel 10-Gigabit Ethernet port as the vPC keepalive interface, you should create a separate VRF for vPC keepalive packets when Layer 3 is enabled with vPC. This process eliminates the possibility of disrupting the vPC keepalive link by the wrong routes learned by a dynamic routing protocol.

This example shows how to configure a new VRF named vpc_keepalive for the vPC keepalive link and how to display the vPC peer keepalive configuration:

```
vrf context vpc_keepalive
interface Ethernet1/31
  switchport access vlan 123
interface Vlan123
  vrf member vpc_keepalive
  ip address 123.1.1.2/30
  no shutdown
vpc domain 1
  peer-keepalive destination 123.1.1.1 source 123.1.1.2 vrf vpc_keepalive

layer3-switch# show vpc peer-keepalive

vPC keep-alive status           : peer is alive
--Peer is alive for             : (154477) seconds, (908) msec
--Send status                   : Success
--Last send at                  : 2011.01.14 19:02:50 100 ms
--Sent on interface             : Vlan123
--Receive status                : Success
--Last receive at               : 2011.01.14 19:02:50 103 ms
--Received on interface         : Vlan123
--Last update from peer         : (0) seconds, (524) msec

vPC Keep-alive parameters
--Destination                   : 123.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc_keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

The services provided by the Cisco Nexus 5500 Platform switch, such as Ping, SSH, Telnet, and RADIUS, are VRF-aware. You must specify the VRF name in the CLI in order to use the correct routing table.

```
layer3-switch# ping 123.1.1.1 vrf vpc_keepalive
PING 123.1.1.1 (123.1.1.1): 56 data bytes
64 bytes from 123.1.1.1: icmp_seq=0 ttl=254 time=3.234 ms
64 bytes from 123.1.1.1: icmp_seq=1 ttl=254 time=4.931 ms
64 bytes from 123.1.1.1: icmp_seq=2 ttl=254 time=4.965 ms
64 bytes from 123.1.1.1: icmp_seq=3 ttl=254 time=4.971 ms
64 bytes from 123.1.1.1: icmp_seq=4 ttl=254 time=4.915 ms

--- 123.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 3.234/4.603/4.971 ms
```

vPC Consistency Check for Layer 3 Parameters

In a vPC topology, vPC peer switches run routing protocols independently and they maintain the routing table independently. Consistency checks are not performed to verify that Layer 3 configurations in the vPC domain are configured symmetrically.

For example, if you configure a router ACL (RACL) on one SVI and you do not configure the router on the corresponding SVI on the vPC peer, a syslog message is not displayed. You must configure the RACL on both devices. This is consistent with the operation of independent routing devices.

Similarly, if you configure peer gateway on one vPC peer and you want the same peer gateway configuration on the other vPC peer, you must configure the peer gateway on the vPC peer.

To confirm that a vPC domain is correctly configured for Layer 3 operations, the following configurations must be consistent:

- SVI configurations
- RACLs
- Routing protocol configurations

Multicast Interaction in a vPC Topology

This section includes the following topics:

Unsupported Multicast Topology

Figure 1-7 shows an unsupported multicast topology in a vPC configuration.

Figure 1-7 Unsupported Multicast Topology with a vPC

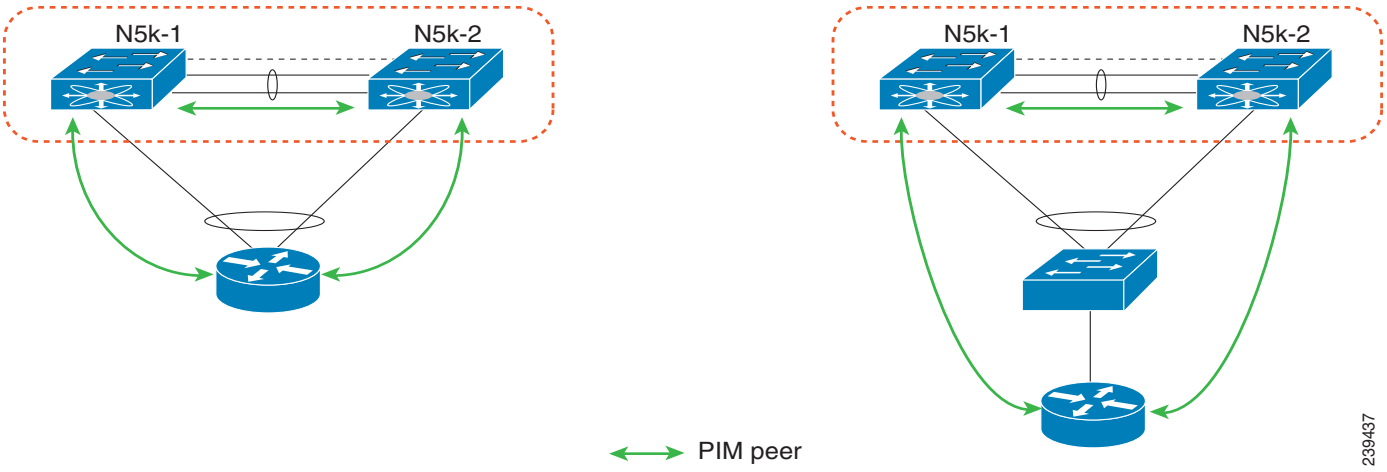

When a PIM router is connected to Cisco Nexus 5500 Platform switches in a vPC topology, the PIM join messages are received only by one switch. The multicast data might be received by the other switch.

Note Multicast forwarding in this topology does not work.

Multicast Routing Table Size

When you enable a vPC on a Nexus 5500 Platform switch, one multicast route (*,G) or (S,G) requires two entries in the routing table; therefore, the multicast routing table size is half the size of what is supported in topologies where vPC is not enabled.

Beginning with Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform multicast routing table size is 2000 entries in non-vPC topologies and 1000 entries in vPC topologies.

Faster Convergence with the Prebuilt Source Tree

In a non-vPC topology, only the designated router (DR) can join the source tree. In a vPC topology, when a receiver is connected to a Cisco Nexus 5500 Platform switch or Fabric Extender (FEX) via vPC, both peer switches initiate a PIM (S,G) join toward the source DR. In a topology where both vPC peer switches have equal costs to the source, the vPC primary switch wins the assert and forwards multicast traffic for receivers connected to the Nexus 5500 Platform switch or FEX using the vPC. The vPC secondary switch also joins the source tree and pulls the multicast data. To prevent data duplication, the vPC secondary switch drops the data due to an empty outgoing interface (OIF) list. Once the vPC secondary switch detects the failure of the vPC primary switch, it adds the receiver VLAN to the OIF list and starts to forward the multicast traffic immediately. Because the vPC secondary switch joins the source tree before the failure, it does not need to initiate the (S,G) join and waits for the tree to be built. As a result, it improves the convergence time in the case of a failure with the active multicast traffic forwarder.

Figure 1-8 shows one receiver that is connected to a dual-homed FEX. The source and Rendezvous Point (RP) are in the Layer 3 network. N5k-2, which is the VPC primary switch, is the multicast traffic forwarder for receivers in VLAN 10.

Figure 1-8 vPC Switch as the Receiver Designated Router

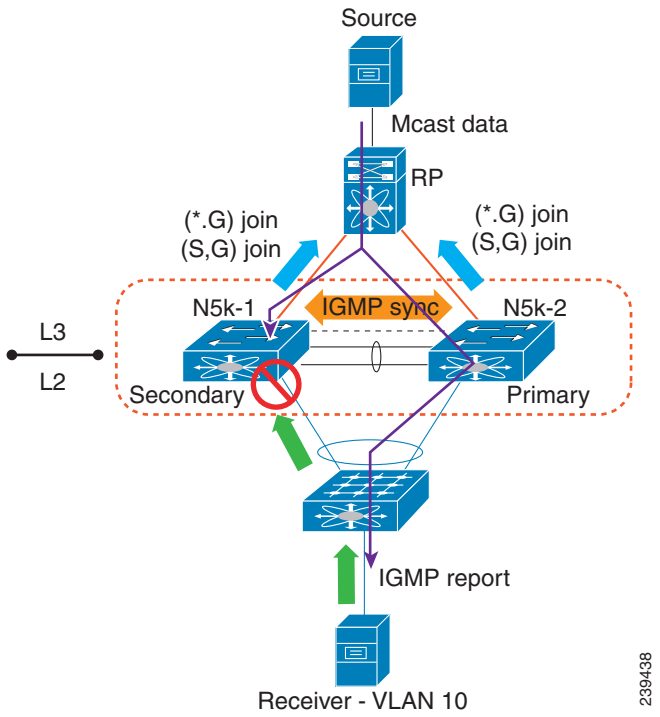

This example shows the output of the multicast routing table and VLAN 10 appears in the OIF list of (S,G) entry on N5k-2. N5k-1 joins the source tree but its OIF list remains empty.

```
N5k-1# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), uptime: 03:03:31, pim ip igmp
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.2
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 03:01:16, igmp

(155.1.3.100/32, 224.1.1.1/32), uptime: 02:13:32, ip pim mrib
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.2
  Outgoing interface list: (count: 0)


N5k-2# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), uptime: 01:48:07, igmp pim ip
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.6
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 01:48:07, igmp

(155.1.3.100/32, 224.1.1.1/32), uptime: 01:00:24, ip pim mrib
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.6
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 00:55:14, mrib
```

The multicast forwarding algorithm applies to all hosts that are connected to the Cisco Nexus 5500 Platform switch or the FEX in a VPC topology, including hosts directly connected to the switch or hosts connected to straight-through FEX topology.

Using a vPC Switch as a Designated Router (PIM DR)

This section includes the following topics:

DR Election and Source Registration

In vPC topologies, a DR election occurs based on the DR priority and the IP address. The elected DR is responsible for sending the source registration toward the RP. When multicast traffic from a directly connected source is received by the non-DR peer switch, the peer switch notifies the DR switch using a Cisco Fabric Services (CFS) message about the source and group address. The DR generates source registration packets to the rendezvous point (RP).

Multicast Data Forwarding

The Cisco Nexus 5500 Platform switch implements a dual-DR mechanism where both vPC peer switches can forward multicast traffic from directly connected sources. The data forwarding rules are as follows:

- The peer switch receives multicast packets from a directly connected source, performs an mroute lookup, and replicates packets for each interface in the OIF list.
- If the OIF is a VLAN trunked over a vPC peer link, one copy is sent over to the peer link for each VLAN that is present in the OIF list. By default, the vPC peer link is considered an mrouter port. Therefore, the multicast packets are sent over to the peer link for each receiving VLAN. You can use the **no ip igmp snooping mrouter vpc-peer link** command to avoid sending multicast traffic over a peer link for each receiver VLAN when there are no orphan ports.

This example shows how to avoid sending the multicast traffic in this scenario:

```
switch-Layer 3-1(config)# no ip igmp snooping mrouter vpc-peer link
Warning: IGMP Snooping mrouter vpc-peer link should be globally disabled on peer VPC
switch as well.
switch-Layer 3-1(config)#
```

With the above CLI configured, the multicast packet is only sent to peer link for VLANs that have orphan ports.

This example shows how to display the list of all orphan ports:

```
switch-Layer 3-1# show vpc orphan-ports
Note:
--------::Going through port database. Please be patient.::--------

VLAN            Orphan Ports
-------         -------------------------
1               Eth1/15
switch-Layer 3-1#
```

Note As of Cisco NX-OS Release 5.0(3)N1(1b), the **no ip igmp snooping mrouter vpc-peer link** command cannot be applied with FEX dual-homed topologies due to a software limitation. The command is used only for interfaces on a Cisco Nexus 5500 Platform switch. This software limitation will be removed in a future software release.

One post-routed multicast packet is sent to a vPC peer link using a reserved VLAN. To configure the reserved VLAN, use the follow commands:

```
switch-Layer 3-1(config)# vpc bind-vrf vrf name vlan VLAN ID
switch-Layer 3-1(config)# vpc bind-vrf default vlan 3000
```

One reserved VLAN is required for each VRF. Without these commands, the receivers in non-vPC VLAN and the receivers connected to Layer 3 interfaces may not receive multicast traffic. The non-vPC VLANs are the VLANs that are not trunked over a peer link.

Multicast traffic that is received over a peer link (with a VLAN ID other than the reserved VLAN ID) is not routed. The multicast traffic is treated as Layer 2 frames that are sent to orphan ports only and not to vPC member ports. The multicast traffic that is received over a peer link with a reserved VLAN ID is routed to a non-vPC VLAN (shown as VLAN 13 in Figure 1-9) and receivers behind the Layer 3 interface. The receivers behind the Layer 3 interface can be hosts directly connected to the Cisco Nexus 5500 Platform switch using Layer 3 interfaces or a router joins the source tree.

Figure 1-9 shows the multicast forwarding rules in a vPC dual-DR topology. In this topology, the source in VLAN 10 and receivers in VLAN 11 and VLAN 12 are the vPC hosts (although in this example they are hosts behind a dual-homed FEX topology where the same rule applies to hosts directly to a Cisco Nexus 5500 Platform switch in a vPC topology). VLAN 13 is a non-vPC VLAN and resides only on N5k-2.

Figure 1-9 Multicast Data Forwarding

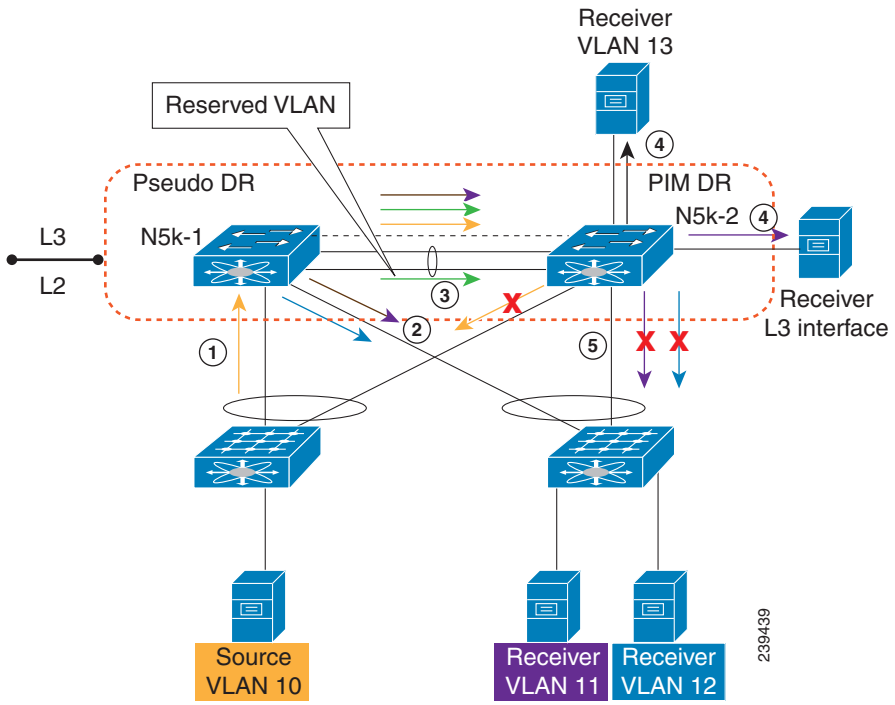

The forwarding process is as follows:

1. IGMP joins from the hosts are synchronized between the two vPC peer switches. N5k-2 is elected as the PIM DR for VLAN 10. Multicast traffic is sent over to N5k-1.
2. The routing engine of N5k-1 performs an mroute lookup and replicates packets to VLAN 11 and VLAN 12. The data packets for VLAN 11 and VLAN 12 are sent to the FEX which in turn sends packets to the two receivers;
3. By default, the replicated packets are sent to the vPC peer link for the source VLAN as well as each receiver VLAN (VLAN 10, VLAN 11, and VLAN 12) in this example. When you use the **no ip igmp snooping mrouter vpc-peer-link** command, the multicast packets are not sent to the peer link for VLAN 10, VLAN 11, and VLAN 12 because there are no orphan ports. One copy of the packets is sent to the peer link with the reserved VLAN 3000 which was configured using the **vpc bind-vrf default vlan 3000** command.

Note In Cisco NX-OS Release 5.0(3)N1(1b), the **no ip igmp snooping mrouter vpc-peer-link** command cannot be applied with a FEX dual-homed topology.

4. For the multicast traffic received from the peer link, if the VLAN ID is the reserved VLAN ID 3000, the N5k-2 route engine performs a Layer 3 lookup and replicates packets to VLAN 13 (a non-vPC VLAN) and receivers behind Layer 3 interfaces.
5. For the multicast packets received over the peer link, VLAN 10, VLAN 11, and VLAN 12 are dropped by N5k-2 to prevent duplicated packets being sent to the vPC hosts. If any orphan ports are in VLAN 10, VLAN 11, and VLAN 12, the packets are bridged to the orphan ports.

Nonfunctional Topologies with Layer 3 and vPC Combined

Some Cisco Nexus 5500 Series switch topologies do not work properly when both Layer 3 and vPC are enabled.

vPC Domain With Layer 3 Enabled on Only One Switch

When two Cisco Nexus 5548/5596 switches are deployed in a vPC domain, both of the switches need to have the same Layer 3 capabilities and the same Layer 3 configuration. The general rule for vPC is that the two devices participating in the vPC domain must have the same functions and capabilities.

Figure 1-10 shows an example of a nonfunctional topology where Layer 3 is enabled on only one switch. When the host sends Layer 3 traffic with the N5k-1 switch's MAC address as the destination MAC address, the traffic could be hashed to the N5k-2 switch. To prevent packet duplication, the traffic received from the peer link is not routed because the assumption is that the peer switch should have routed the traffic and all of the traffic received from the peer link should only be bridged.

Figure 1-10 Nonfunctional Topology: Layer 3 Enabled on Only One vPC Switch

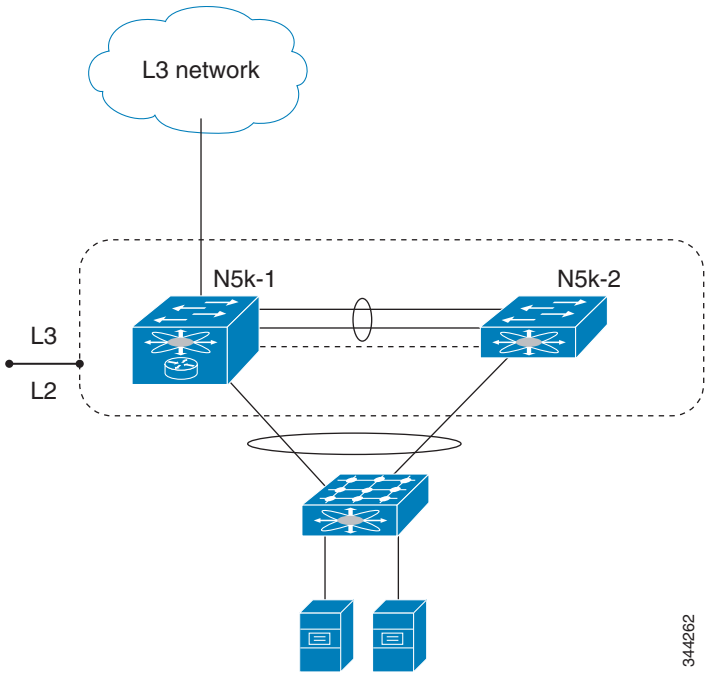

This mismatched Layer 3 configuration can happen in the following scenarios:

- Only one Cisco Nexus 5000 series switch has a Layer 3 module or only one Cisco Nexus 5000 series switch has the Layer 3 license installed.
- Both Cisco Nexus 5000 switches have a Layer 3 module and the Layer 3 license installed, but only one Cisco Nexus 5000 series switch has SVI configured.
- Both Cisco Nexus 5000 Series switches have an SVI configured, but only one Cisco Nexus 5000 Series switch has the First Hop Redundancy Protocol (FHRP) configured.

In all these scenarios, the traffic forwarding does not work properly. Additionally, we recommend that you have identical configurations for all other Layer 3 parameters, such as Router ACLs (RACLs) and routing protocols.

Layer 3 parameters are not part of the vPC consistency check. So, you have to manually verify that the Layer 3 configurations are identical on both Cisco Nexus 5000 Series switches.

Topology with an Additional Parallel Link Between Two Switches

Figure 1-11 shows the nonfunctional topology where a parallel link in addition to the vPC peer link and the vPC peer-keepalive link between the two switches, and the two switches have Layer 3 enabled. You can have a link between the two switches using the front panel ports for the vPC peer-keepalive link, but this link should only be used to carry vPC-keepalive message.

In some circumstances, you might consider having a separate link between the two vPC switches, either to carry non-vPC VLAN traffic or to form Layer 3 routing protocol peering. While this design is supported on the Cisco Nexus 7000 Series switch, it does not work on the Cisco Nexus 5000 Series switch. With the Cisco Nexus 5000 Series switch, we recommend that you use a vPC peer link for Layer 3 peering to carry both vPC and non-vPC VLAN traffic.

Figure 1-11 Nonfunctional Topology: Additional Parallel Peer Link Between Switches

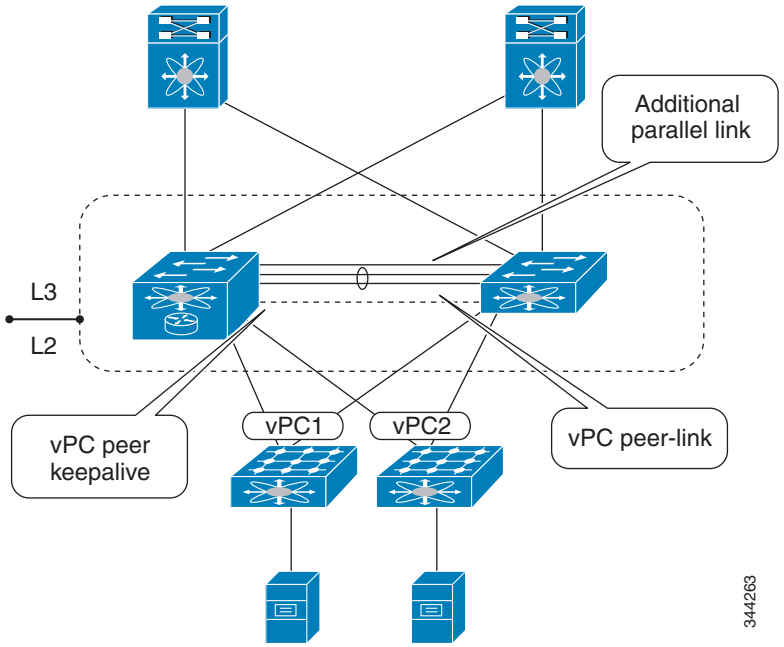

Connecting a Router Using a VLAN Trunk Port

Figure 1-12 shows the nonfunctional topology where a router is connected to a Cisco Nexus 5000 Series switch using a VLAN trunk port and PIM is enabled for the same VLAN.

Figure 1-12 Nonfunctional Topology: Connecting a Router Using a VLAN Trunk Port

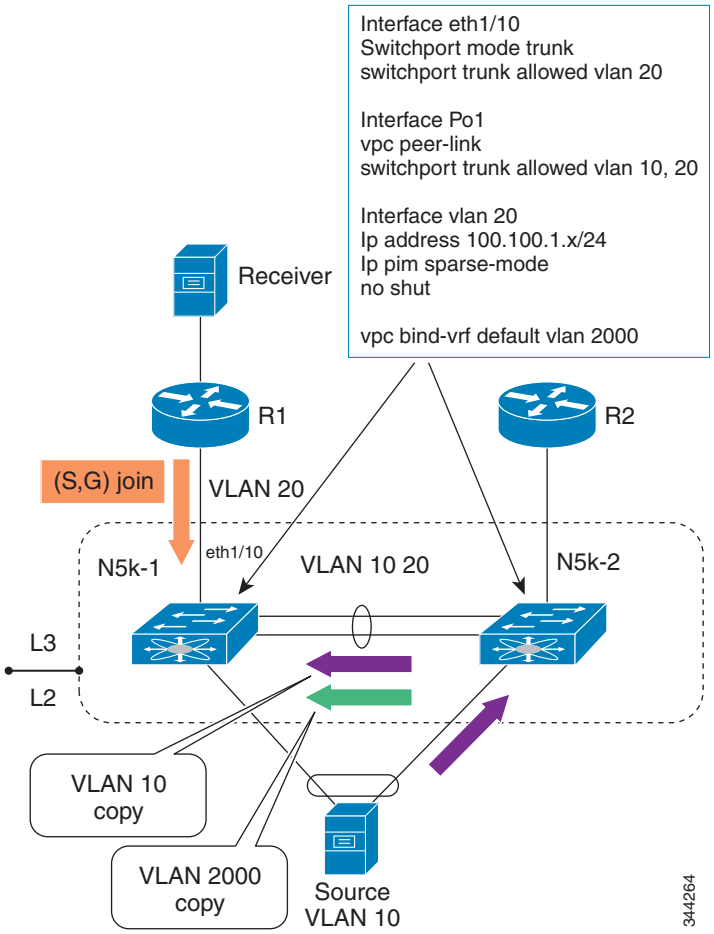

The R1 and R2 routers can be any platform that support both the Layer 2 and Layer 3 functions, such as a Cisco Nexus 7000 Series switch, a Cisco Catalyst 6500 Series switch, or a Cisco Catalyst 4900 Series switch. The intention of this design is to extend a VLAN to all four devices.

From the Layer 3 point of view, the topology has four devices in same VLAN. Let us assume that the PIM (S,G) join message is sent from the R1 router to the N5k-1 switch. It is possible that a source behind the vPC can send traffic to the N5k-2 switch. The N5k-2 switch sends two copies of the multicast packets to the peer link. One copy is for source VLAN 10, and the second copy is for the special VLAN 2000 that is configured with the **vpc bind-vrf default vlan** command. When the N5k-1 switch receives the packet through the peer link for VLAN 10, it only conducts the Layer 2 bridging. That is, it only sends the packet to the orphan ports that reside in VLAN 10. In addition, the N5k-1 switch tries to route the multicast packet for the VLAN 2000 copy that it received from the peer link. In order to prevent packet duplication, the N5k-1 swtich only routes the multicast packet to the Layer 3 interface or the non-vPC VLAN. (In this example, VLAN 20 is trunked over the peer-link and is considered to be a vPC VLAN.) Therefore, the N5k-1 switch does not route the multicast packet to the R1 router.

In this scenario, if the intention is to extend the VLAN to all four devices, the alternative design is to not enable Layer 3 on the N5k-1 and N5k-2 switches. Such topology is supported if the vPC is replaced with vPC+, which requires FabricPath. For more information, see the *Cisco Nexus 5000 Series NX-OS FabricPath Operations Guide, Release 5.1(3)N1(1)*.

Routing Peering Over vPC

Figure 1-13 shows a nonfunctional topology where the dynamic routing protocol is enabled between a router and two Cisco Nexus 5000 Series switches in same vPC domain. The PIM protocol does not work in this topology design, and while the unicast routing protocol allows peering in the vPC, we do not recommend this design. When Enhanced vPC is deployed between a pair of Cisco Nexus 5000 Series switches, the routing peering topology shown in Figure 1-13 is supported. Enhanced vPC requires FabricPath. For more informaiton, see the *Cisco Nexus 5000 Series NX-OS FabricPath Operations Guide, Release 5.1(3)N1(1).*

Figure 1-13 Nonfunctional Topology: Routing Peering Over vPC

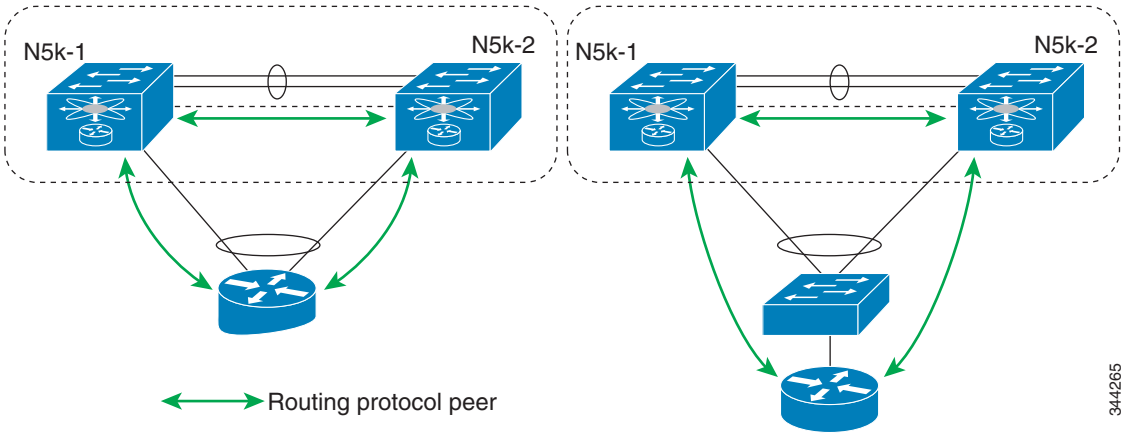

Software Upgrade and Downgrade Impact

In Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform switch does not support ISSUs when Layer 3 modules are installed and Layer 3 features are enabled. Use the **install all** command and the **show install all impact** command to determine the impact of the software upgrade and to indicate whether the software upgrade with Layer 3 features enabled will be disruptive and would require a switch and FEX reload.

show install all impact kickstart

This example shows the output of the **show install all** command:

```
Layer 3-N5548-2# show install all impact kickstart
n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg system n5000-uk9.5.0.3.N1.0.271.bin.upg

Verifying image bootflash:/n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg for boot variable
"kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[############        ]  50%
[####################] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:/n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "fexth" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes      disruptive         reset  Non-disruptive install not supported if
Layer 3 was enabled
   100       yes      disruptive         reset  Non-disruptive install not supported if
Layer 3 was enabled



Images will be upgraded according to following table:
Module       Image         Running-Version             New-Version  Upg-Required
------  ----------  ----------------------  ----------------------  ------------
     1      system             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1   kickstart             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1        bios      v3.4.0(01/13/2011)      v3.4.0(01/13/2011)           no
   100       fexth             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1   power-seq                    v3.0                   v3.0           no
     2   power-seq                    v1.0                   v1.0           no
     1          uC               v1.0.0.14              v1.0.0.14           no

Layer 3-N5548-2#
```

You can perform a nondisruptive ISSU from an earlier release to NX-OS Release 5.0(3)N1(1b) when upgrading without Layer 3 features enabled.

show spanning-tree issu-impact

To verify that the current STP topology is consistent with ISSU requirements, use the **show spanning-tree issu-impact** command to display the STP configuration and whether or not there are potential STP issues.

This example shows how to display information about the STP impact when performing an ISSU:

```
nexus5010# show spanning-tree issu-impact
For ISSU to Proceed, Check the Following Criteria :
1. No Topology change must be active in any STP instance
2. Bridge assurance(BA) should not be active on any port (except MCT)
3. There should not be any Non Edge Designated Forwarding port (except MCT)
4. ISSU criteria must be met on the VPC Peer Switch as well
```

```
Following are the statistics on this switch

No Active Topology change Found!
Criteria 1 PASSED !!

No Ports with BA Enabled Found!
Criteria 2 PASSED!!

No Non-Edge Designated Forwarding Ports Found!
Criteria 3 PASSED !!

ISSU Can Proceed! Check Peer Switch.
```

For information on upgrade procedures, see the *Cisco Nexus 5000 Series NX-OS Upgrade and Downgrade Guide*.