CHAPTER 3

# Understanding Network-Level High Availability

This chapter describes Cisco NX-OS network high availability and includes the following sections:

# Information About Network-Level High Availability

Cisco NX-OS network-level HA is optimized by tools and functionality that provide failovers and fallbacks transparently and quickly. The features described in this chapter ensure high availability at the network level.

## Virtualization Support

Each virtual device context (VDC) in a system runs a separate Spanning Tree Protocol (STP), which includes extensions to support virtualization. Each VDC can also run one or more instances of a routing protocol. The network-level HA features described in this chapter apply to a failure or restart of a VDC in the same manner as a failure or restart of the system.

**Note** For complete information on VDCs, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.1*.

# Licensing Requirements

The following table shows the licensing requirements for network-level high availability features:

| Product | License Requirement |
|---|---|
| Cisco NX-OS | The network-level high availability features require no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided for free. |
| VDC | VDC requires an Advanced Services license. |
| BGP | Border Gateway Protocol (BGP) requires an Enterprise Services license. |

For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

# Spanning Tree Protocol

**Note** Spanning Tree Protocol (STP) is used to refer to IEEE 802.1w and IEEE 802.1s. If this publication is referring to the IEEE 802.1D STP, 802.1D is stated specifically.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. Multiple active paths between end stations cause loops in the network that result in network devices learning end station MAC addresses on multiple Layer 2 LAN ports. This condition can result in a broadcast storm, which creates an unstable network.

STP provides a loop-free network at the Layer 2 level. Layer 2 LAN ports send and receive STP frames, which are called Bridge Protocol Data Units (BPDUs), at regular intervals. Network devices do not forward these frames but use the frames to determine the network topology and to construct a loop-free path within that topology. Using the spanning tree topology, STP forces redundant data paths into a blocked state. If a network segment in the spanning tree fails and a redundant path exists, the STP algorithm recalculates the spanning tree topology and activates the blocked path.

Cisco NX-OS also supports Multiple Spanning Tree Protocol (MSTP). The multiple independent spanning tree topology enabled by MSTP provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of STP instances required to support a large number of VLANs.

MST incorporates Rapid Spanning Tree Protocol (RSTP), which allows rapid convergence. MST improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

**Note** You can configure spanning tree parameters only on Layer 2 interfaces; a spanning tree configuration is not allowed on a Layer 3 interface. For information on creating Layer 2 interfaces, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x*.

For details about STP behavior and configuration, see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*.

# Virtual Port Channels

The major limitation in classic port channel communication is that the port channel operates only between two devices. In large networks, the support of multiple devices together is often a design requirement to provide some form of hardware failure alternate path. This alternate path is often connected in a way that would cause a loop, limiting the benefits gained with port channel technology to a single path. To address this limitation, Cisco NX-OS provides a technology called virtual port channel (vPC). Although a pair of switches acting as a vPC peer endpoint looks like a single logical entity to port channel-attached devices, the two devices that act as the logical port channel endpoint are still two separate devices. This environment combines the benefits of hardware redundancy with the benefits of port channel loop management.

For more information on vPCs, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x*.

# First-Hop Redundancy Protocols

Within a group of two or more routers, first-hop redundancy protocols (FHRPs) allow a transparent failover of the first-hop IP router. Cisco NX-OS supports the following FHRPs:

- Hot Standby Router Protocol (HSRP)—HSRP provides first-hop routing redundancy for IP hosts on Ethernet networks configured with a default gateway IP address. An HSRP router group of two or more routers chooses an active gateway and a standby gateway. The active gateway routes packets while the standby gateway remains idle until the active gateway fails or when preset conditions are met.

  Many host implementations do not support any dynamic router discovery mechanisms but can be configured with a default router. Running a dynamic router discovery mechanism on every host is not feasible for a number of reasons, including administrative overhead, processing overhead, and security issues. HSRP provides failover services to these hosts.

- Virtual Router Redundancy Protocol (VRRP)—VRRP dynamically assigns responsibility for one or more virtual routers to the VRRP routers on a LAN, which allows several routers on a multi-access link to use the same virtual IP address. A VRRP router is configured to run VRRP with one or more other routers attached to a LAN. One router is elected as the virtual router master, while the other routers act as backups if the virtual router master fails.

- Gateway Load Balancing Protocol (GLBP)—GLBP provides path redundancy for IP by sharing protocol and Media Access Control (MAC) addresses between redundant gateways. In addition, GLBP allows a group of Layer 3 routers to share the load of the default gateway on a LAN. A GLBP router can automatically assume the forwarding function of another router in the group if the other router fails.

  GLBP performs a similar function to HSRP and the Virtual Router Redundancy Protocol (VRRP), which allows multiple routers to participate in a virtual group configured with a virtual IP address. GLBP performs an additional load balancing function that HSRP and VRRP do not provide. GLBP shares the forwarding load among all routers in a GLBP group instead of allowing a single router to handle the entire load while the other routers remain idle. HSRP and VRRP elect one member as the active router to forward packets to the virtual IP address for the group. The other routers in the group are redundant until the active router fails.

For configuration details about FHRPs, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x*.

# Nonstop Forwarding in Routing Protocols

The Nexus 7000 series supports Open Shortest Path First version 2 (OSPFv2), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP), each of which includes network-level HA mechanisms to minimize network disruption because of process restarts or supervisor switchovers.

As an example, the HA features of OSPFv2 are described in this section. For HA configuration details on OSPFv2, EIGRP, and BGP, see the “High Availability and Graceful Restart” chapter of the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x*.

This section includes the following topics:

## OSPFv2 Stateless Restart

If a Cisco NX-OS system that runs OSPFv2 experiences a cold reboot, the network stops forwarding traffic to the system and removes the system from the network topology. In this scenario, OSPFv2 experiences a stateless restart and removes all neighbor adjacencies on the local system. Cisco NX-OS applies the startup configuration and OSPFv2 rediscovers the neighbors and establishes the adjacencies again.

## OSPFv2 Graceful Restart on a Switchover

When a supervisor switchover begins, OSPFv2 initiates a graceful restart, or nonstop forwarding (NSF), by announcing that OSPFv2 will be unavailable for some time. During the switchover, neighbor devices continue to forward traffic and keep the system in the network topology. After the switchover, Cisco NX-OS applies the running configuration, and OSPFv2 informs the neighbors that it is operational again. The neighbor devices help to reestablish adjacencies.

## OSPFv2 Graceful Restart on an OSFPv2 Process Failure

OSPFv2 automatically restarts if the process experiences problems. After the restart, OSPFv2 initiates a graceful restart so that the platform is not taken out of the network topology. If you manually restart OSPF, it performs a graceful restart, which is similar to a stateful switchover. The running configuration is applied in both cases. The graceful restart allows OSPFv2 to remain in the data forwarding path through the process restart.

**Note** If the restarting OSPFv2 interface does not come back up before the end of the grace period, or if the network experiences a topology change, the OSPFv2 neighbors tear down adjacency with the restarting OSPFv2 and treat it as a normal OSPFv2 restart.

# Additional References

For additional information related to implementing network-level HA features, see the following sections:

## Related Documents

| Related Topic | Document Title |
|---|---|
| Virtual device context (VDC) | *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.1* |
| Graceful restart | "High Availability and Graceful Restart" chapter of the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x* |
| In-service software upgrades (ISSU) | Chapter 5, "Understanding In-Service Software Upgrades" |
| Licensing | *Cisco NX-OS Licensing Guide* |

## Standards

| Standards | Title |
|---|---|
| IEEE 802.1Q-2006 (formerly known as IEEE 802.1s), IEEE 802.1D-2004 (formerly known as IEEE 802.1w), IEEE 802.1D, IEEE 802.1t | — |

## MIBs

| MIBs | MIBs Link |
|---|---|
| • CISCO-SYSTEM-EXT-MIB: ciscoHaGroup, cseSwCoresTable, cseHaRestartNotify, cseShutDownNotify, cseFailSwCoreNotify, cseFailSwCoreNotifyExtended<br>• CISCO-STP-EXTENSION-MIB<br>• CISCO-PROCESS-MIB<br>• CISCO-RF-MIB | To locate and download MIBs, go to the following URL:<br>http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml |

## RFCs

| RFCs | Title |
|---|---|
| No RFCs are supported by this feature | — |

# Technical Assistance

| Description | Link |
|---|---|
| Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content. | http://www.cisco.com/cisco/web/support/index.html |