



Network-Level High Availability

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

- [About Network-Level High Availability, on page 1](#)
- [Spanning Tree Protocol, on page 1](#)
- [Virtual Port Channels, on page 2](#)
- [First-Hop Redundancy Protocols, on page 2](#)
- [Nonstop Forwarding in Routing Protocols, on page 3](#)
- [Additional References for Network-Level High Availability, on page 3](#)

About Network-Level High Availability

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.

Spanning Tree Protocol



Note The Spanning Tree Protocol (STP) refers 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 the 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 the 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 9000 Series NX-OS Interfaces Configuration Guide](#).

For details about STP behavior and configuration, see the [Cisco Nexus 9000 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 9000 Series NX-OS Interfaces Configuration Guide](#).

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 primary virtual router, while the other routers act as backups if the primary virtual router fails.

For configuration details about FHRPs, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).

Nonstop Forwarding in Routing Protocols

Cisco NX-OS provides a multilevel high-availability architecture. Open Shortest Path First version 2 (OSPFv2) supports stateful restart, which is also referred to as nonstop routing (NSR). If OSPFv2 experiences problems, it attempts to restart from its previous runtime state. The neighbors would not register any neighbor event in this case.

If the first restart is not successful and another problem occurs, OSPFv2 attempts a graceful restart. A graceful restart, or nonstop forwarding (NSF), allows OSPFv2 to remain in the data forwarding path through a process restart. When OSPFv2 needs to do a graceful restart, it first sends a link-local opaque (type 9) link-state advertisement (LSA), called a grace LSA. (For more information about opaque LSAs, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).) The restarting of the OSPFv2 platform is called NSF capable. The grace LSA includes a grace period, which is a specified time that the neighbor OSPFv2 interfaces hold onto the LSAs from the restarting OSPFv2 interface. (Typically, OSPFv2 tears down the adjacency and discards all LSAs from a down or restarting OSPFv2 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv2 interface as if the interface were still adjacent. When the restarting OSPFv2 interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its LSA updates again. At this point, the NSF helpers recognize that graceful restart has finished.

Scenarios where a stateful restart is used:

- First recovery attempt after a process experiences problems.
- User-initiated switchover using the **system switchover** command.
- Active supervisor removal.
- Active supervisor reload using the **reload module active-sup** command.

Scenarios where graceful restart is used:

- Second recovery attempt after a process experiences problems within a 4-minute interval.
- Manual restart of the process using the **restart {ospfv3 | ospf}** command.



Note For more information on nonstop routing in routing protocols, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).

Additional References for Network-Level High Availability

This section describes additional information related to network-level high availability.

Related Documents

Related Topic	Document Title
Graceful restart	Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration
Licensing	Cisco NX-OS Licensing Guide

MIBs

MIBs	MIBs Link
MIBs related to network-level high availability	For more information about MIBs and to download MIBs from the MIBs link, refer to Cisco Nexus 7000 Series and 9000 Series MIB Quick Reference .