



# CHAPTER 1

## Overview

---

Cisco NX-OS is a resilient operating system that is specifically designed for high availability at the network, system, and process level.

This chapter describes high availability (HA) concepts and features for Cisco NX-OS devices and includes the following sections:

- [Information About High Availability, page 1-1](#)
- [Service-Level High Availability, page 1-2](#)
- [System-Level High Availability, page 1-2](#)
- [Network-Level High Availability, page 1-4](#)
- [Additional Management Tools for Availability, page 1-5](#)

## Information About High Availability

To prevent or minimize traffic disruption during hardware or software failures, Cisco NX-OS has these features:

- **Redundancy**—Cisco NX-OS HA provides physical and software redundancy at every component level, spanning across the physical, environmental, power, and system software aspects of its architecture.
- **Isolation of planes and processes**—Cisco NX-OS HA provides isolation between control and data forwarding planes within the device and between software components, so that a failure within one plane or process does not disrupt others.
- **Restartability**—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- **Supervisor stateful switchover**—The Nexus 7000 series supports an active/standby dual supervisor configuration. State and configuration remain constantly synchronized between the two supervisor modules to provide seamless and stateful switchover in the event of a supervisor module failure.
- **Nondisruptive upgrades**—Cisco NX-OS supports the in-service software upgrade (ISSU) feature, which allows you to upgrade the device software while the switch continues to forward traffic. ISSU reduces or eliminates the downtime typically caused by software upgrades.

[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)

## Service-Level High Availability

Cisco NX-OS has a modularized architecture that compartmentalizes components for fault isolation, redundancy, and resource efficiency.

For additional details about service-level HA, see [Chapter 2, “Understanding Service-Level High Availability.”](#)

This section includes the following topics:

- [Isolation of Processes, page 1-2](#)
- [Process Restartability, page 1-2](#)

### Isolation of Processes

In the Cisco NX-OS software, independent processes, known as *services*, perform a function or set of functions for a subsystem or feature set. Each service and service instance runs as an independent, protected process. This approach provides a highly fault-tolerant software infrastructure and fault isolation between services. A failure in a service instance (such as 802.1Q) will not affect any other services running at that time (such as the Link Aggregation Control Protocol[LACP]). Additionally, each instance of a service can run as an independent process, which means that two instances of a routing protocol (for example, two instances of the Open Shortest Path First [OSPF] protocol) can run as separate processes.

### Process Restartability

Cisco NX-OS processes run in a protected memory space independently from each other and the kernel. This process isolation provides fault containment and enables rapid restarts. Process restartability ensures that process-level failures do not cause system-level failures. In addition, most services can perform stateful restarts, which allows a service that experiences a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network.

## System-Level High Availability

The Nexus 7000 series is protected from system failure by redundant hardware components and a high-availability software framework.

For additional information about system-level HA features, see [Chapter 3, “Understanding System-Level High Availability.”](#)

This section includes the following topics:

- [Physical Redundancy, page 1-3](#)
- [ISSU, page 1-3](#)
- [VDCs, page 1-3](#)

**[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)**

## Physical Redundancy

The Nexus 7000 series has the following physical redundancies:

- **Power Supply Redundancy**—The Cisco Nexus 7000 series chassis supports three power supply modules, each of which is composed of two internalized isolated power units, giving it two power paths per modular power supply, and six paths in total, per chassis, when fully populated.
- **Fan Tray Redundancy**—The Cisco Nexus 7000 series chassis contains two redundant system fan trays for I/O module cooling and two additional fan trays for switch fabric module cooling. Only one of each pair of fan trays is sufficient to provide system cooling. In the case of a fan tray failure, you must leave the failed unit in place to ensure proper airflow until a replacement is available. The fan trays are hot swappable, but you must complete the removal and replacement within three minutes to avoid an automatic system shutdown.
- **Fabric Redundancy**—Cisco NX-OS provides switching fabric availability through redundant switch fabric modules. You can configure a single Cisco Nexus 7000 series chassis with one to five switch fabric cards for capacity and redundancy. Each I/O module installed in the system automatically connects to and uses all functionally installed switch fabric modules. A failure of a switch fabric module triggers an automatic reallocation and balancing of traffic across the remaining active switch fabric modules. Replacing the failed fabric module reverses this process. Once you insert the fabric module and bring it online, traffic is again redistributed across all installed fabric modules and redundancy is restored.
- **Supervisor Module Redundancy**—The Cisco Nexus 7000 series chassis supports dual supervisor modules to provide redundancy for the control and management plane. A dual supervisor configuration operates in an active/standby capacity in which only one of the supervisor modules is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two supervisor modules to provide a stateful switchover if the active supervisor module fails.

For additional details about physical redundancy in the Nexus 7000 series, see [Chapter 3, “Understanding System-Level High Availability.”](#)

## ISSU

Cisco NX-OS allows you to perform an in-service software upgrade (ISSU), which is also known as a nondisruptive upgrade. The modular software architecture of NX-OS supports plug-in-based services and features, which allow you to perform complete image upgrades of supervisors and switching modules with little to no impact on other modules. Due to this design, you can upgrade NX-OS nondisruptively with no impact to the data forwarding plane and allow for nonstop forwarding during a software upgrade, even between full image versions.

For additional details about ISSU, see [Chapter 5, “Understanding In-Service Software Upgrades.”](#)

## VDCs

Cisco NX-OS implements a logical virtualization at the device level, which allows multiple instances of a device to operate on the same physical switch simultaneously. These logical operating environments are known as *virtual device contexts*, or VDCs. VDCs provide logically separate device environments that you can independently configure and manage. This degree of isolation provides fault isolation in addition to security and administrative benefits. Human error or failure conditions due to the

**[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)**

configuration are isolated within a given virtual device. While virtual device contexts are not primarily a high-availability feature, the operationally independent fault domains contribute to availability and prevent service disruptions that are associated with device configuration.

For more information on VDCs, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0*.

## Network-Level High Availability

Network convergence is optimized by providing tools and functions to make both failover and fallback transparent and fast.

For additional information about network-level HA features, see [Chapter 4, “Understanding Network-Level High Availability.”](#)

This section includes the following topics:

- [Layer 2 HA Features, page 1-4](#)
- [Layer 3 HA Features, page 1-4](#)

### Layer 2 HA Features

Cisco NX-OS provides these Layer 2 HA features:

- Spanning Tree Protocol enhancements, such as bridge protocol data unit (BPDU) Guard, Loop Guard, Root Guard, BPDU Filters, and Bridge Assurance, to guarantee the health of the Spanning Tree Protocol control plane
- Unidirectional Link Detection (UDLD) Protocol
- Shortest Path First (SPF) optimizations such as link-state advertisement (LSA) pacing and incremental SPF
- IEEE 802.3ad link aggregation

### Layer 3 HA Features

Cisco NX-OS provides these Layer 3 HA features:

- Nonstop forwarding (NSF) graceful restart extensions for routing protocols  
OSPFv2, OSPFv3, Intermediate System to Intermediate System (IS-IS), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) utilize graceful restart extensions to the base protocols to provide nonstop forwarding and least obtrusive routing recovery for those environments.
- Protocol-based periodic refresh
- Millisecond timers for First-Hop Redundancy Protocols (FHRP) such as Hot Standby Router Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), and Virtual Router Redundancy Protocol (VRRP)

**[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)**

## Additional Management Tools for Availability

Cisco NX-OS incorporates several Cisco system management tools for monitoring and notification of system availability events.

This section includes the following topics:

- [GOLD, page 1-5](#)
- [EEM, page 1-5](#)
- [Smart Call Home, page 1-5](#)

### GOLD

Cisco Generic On-Line Diagnostics (GOLD) subsystem and additional monitoring processes on the supervisor facilitate the triggering of a stateful failover to the redundant supervisor upon the detection of unrecoverable critical failures, service restartability errors, kernel errors, or hardware failures.

For information about configuring GOLD, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

### EEM

Cisco Embedded Event Manager (EEM) consists of Event Detectors, the Event Manager, and an Event Manager Policy Engine. Using EEM, you can define policies to take specific actions when the system software recognizes certain events through the Event Detectors. The result is a flexible set of tools to automate many network management tasks and to direct the operation of Cisco NX-OS to increase availability, collect information, and notify external systems or personnel about critical events.

For information about configuring EEM, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

### Smart Call Home

Combining Cisco GOLD and Cisco EEM capabilities, Smart Call Home provides an e-mail-based notification of critical system events. Smart Call Home has message formats that are compatible with pager services, standard e-mail, or XML-based automated parsing applications. You can use this feature to page a network support engineer, e-mail a network operations center, or use Cisco Smart Call Home services to automatically generate a case with Cisco's Technical Assistance Center (TAC).

For information about configuring Smart Call Home, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

***Send document comments to [nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)***