CHAPTER 1

Overview

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

Information About High Availability

The purpose of High Availability (HA) is to limit the impact of failures—both hardware and software—within a system. The Cisco NX-OS operating system is designed for high availability at the network, system, and service levels.

The following Cisco NX-OS features minimize or prevent traffic disruption in the event of a failure:

- Redundancy— redundancy at every aspect of the software architecture.
- Isolation of processes— isolation between software components to prevent a failure within one process disrupting other processes.
- 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— Active/standby dual supervisor configuration. State and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide seamless and stateful switchover in the event of a VSM failure.

The Cisco Nexus 1000V system is made up of the following:

- Virtual Ethernet Modules (VEMs) running within virtualization servers. These are represented as modules within the VSM.
- A remote management component, for example. VMware vCenter Server.
- One or two VSMs running within Virtual Machines (VMs).

Figure 1-1 shows the HA components and the communication links between them.

Figure 1-1 Cisco Nexus 1000V HA Components and Communication Links

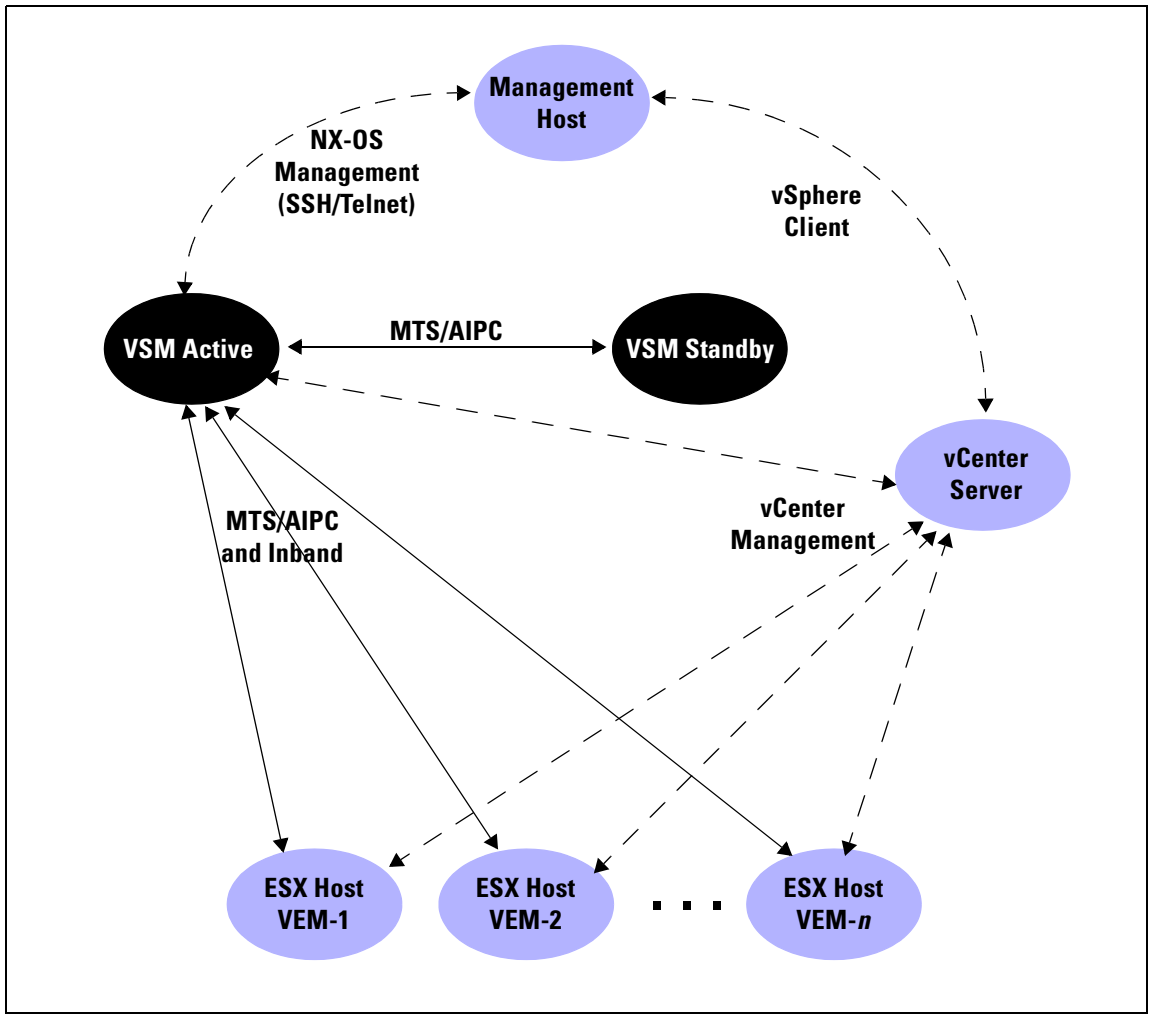

Service-Level High Availability

The Cisco NX-OS software compartmentalizes processes for fault isolation, redundancy, and efficiency.

This section includes the following topics:

- Isolation of Processes, page 1-3
- Process Restartability, page 1-3

For additional details about service-level HA, see Chapter 2, "Understanding Service-Level High Availability."

Isolation of Processes

Cisco NX-OS software has independent processes, known as *services,* that 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 provides a highly fault-tolerant software infrastructure and fault isolation between services. A failure in a service instance will not affect any other services running at that time. Additionally, each instance of a service can run as an independent process, which means that two instances of a routing protocol can run as separate processes.

Process Restartability

Cisco NX-OS processes run in a protected memory space independently of 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 Cisco Nexus 1000V supports redundant VSM virtual machines — a primary and a secondary — running as an HA pair. Dual VSMs operate in an active/standby capacity in which only one of the VSMs is active at any given time, while the other acts as a standby backup. The VSMs are configured as either primary or secondary as a part of the Cisco Nexus 1000V installation. The state and configuration remain constantly synchronized between the two VSMs to provide a stateful switchover if the active VSM fails.

Single or Dual Supervisors

The Cisco Nexus 1000V system is made up of the following:

- Virtual Ethernet Modules (VEMs) running within virtualization servers (these are represented as modules within the VSM).
- A remote management component, for example, VMware vCenter Server.
- One or two Virtual Supervisor Modules (VSMs) running within Virtual Machines (VMs).

For more information about system-level high availability, see the "Configuring System-Level High Availability" section on page 3-1.

Network-Level High Availability

The Cisco Nexus 1000V HA at the network level includes port channels and Link Aggregation Control Protocol (LACP). A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to eight physical links. If a member port within a port channel fails, the traffic previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, LACP lets you configure up to 16 interfaces into a port channel. A maximum of eight interfaces can be active, and a maximum of eight interfaces can be placed in a standby state.

For additional information about port channels and LACP, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(1)*.