



Overview

This chapter contains the following sections:

- [Information About High Availability, page 1](#)
- [System Components, page 1](#)
- [InterCloud Link, page 2](#)
- [Service-Level High Availability, page 2](#)
- [System Level High Availability, page 2](#)
- [InterCloud Extender High Availability Handshake, page 3](#)
- [InterCloud Extender High Availability Heartbeats, page 3](#)
- [Split Brain Resolution, page 4](#)

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.

System Components

The Cisco Nexus 1000V InterCloud consists of the following components:

- Cisco Prime Network Controller
- Cisco Nexus 1000V InterCloud Virtual Supervisor Module (VSMs)
- InterCloud Switch
- InterCloud Extender
- Virtual Ethernet modules (VEMs) that are represented as modules within the VSM
- A remote management component such as VMware vCenter Server.

See the *Cisco Nexus 1000V InterCloud Installation Guide* for more information about the various system components.

InterCloud Link

The Cisco Nexus 1000V InterCloud provides a system level high availability solution via redundant InterCloud Links. An InterCloud Link is a secure connection between an enterprise and public cloud. A secure Layer 2 tunnel connects the InterCloud Extender and InterCloud Switch, thereby extending the enterprise network into the cloud. The InterCloud Link is considered a single unit consisting of the InterCloud Extender in the enterprise and InterCloud Switch in the public cloud.

Service-Level High Availability

Isolation of Processes

The 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 way of operating provides a highly fault-tolerant software infrastructure and fault isolation between services. A failure in a service instance does not affect any other services that are 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. These stateful restarts allow 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 InterCloud 1000V supports redundant InterCloud Links, a primary and a secondary, running as an HA pair. Dual InterCloud Links operate in an active/standby capacity in which only one of the InterCloud

Links is active at any given time, while the other acts as a standby backup. The InterCloud Links are configured as primary or secondary during the Cisco Nexus 1000V InterCloud installation.

Redundancy Manager

Redundancy Manager is the service on the Cisco Nexus 1000V InterCloud that manages the high availability feature within and among gateways to provide a system-level high availability solution. Redundancy Manager within each gateway communicates with its peer gateways to ensure the system is in a healthy state.

InterCloud Extender High Availability Handshake

In the Cisco Nexus 1000V InterCloud HA model, the InterCloud Link consisting of InterCloud Extender in the enterprise and the InterCloud Switch in a provider cloud are together considered a single unit. In an HA deployment, a second, standby InterCloud Link exists to minimize the impact of a failure on the active InterCloud Link.

In HA mode, two InterCloud Links are deployed in the system. During instantiation, the Cisco Prime Network Services Controller assigns an HA role (either primary or secondary) to each InterCloud Extender. Once configured with their local and peer information and the site-to-site tunnel has been established, the InterCloud Extenders perform a handshake over UDP port 9984. Upon initial handshake, the InterCloud Extenders use their HA roles to determine which should move to the active or standby states. Normally, the InterCloud Extender with the Primary HA role will move active and the InterCloud Extender with the Secondary HA role will move standby.

It is possible, however, that a Secondary InterCloud Extender moves active if it cannot communicate with the Primary InterCloud Extender during initial handshake. If this occurs, the Primary will move to the standby state once it has come up and performed the handshake with the Secondary InterCloud Extender.

If the active InterCloud Link fails, the standby InterCloud Link moves to active state and the failed InterCloud Link is rebooted and moves to the standby state.

InterCloud Extender High Availability Heartbeats

After the handshake has occurred between the primary and secondary InterCloud Extenders, they exchange heartbeats to share data and ensure the system is healthy. Similar to the handshake, the heartbeats are sent and received on UDP port 9984.

The heartbeats include useful information such as local and peer states, control flags, and tunnel status that allow the Redundancy Manager on each InterCloud Extender to make intelligent decisions regarding the health of the system as a whole.

The following intervals apply when sending heartbeat messages.

Interval	Description
5 seconds	Interval at which heartbeat requests are sent.
35 seconds	Interval after which missed heartbeats indicate degraded communication on the management interface so that heartbeats are also sent through the site-to-site secure tunnel and InterCloud Switches.

Interval	Description
300 seconds	Interval after which the standby InterCloud Link will reload should a WAN connectivity issue occur. This protects against the possibility of a failure occurring on both InterCloud Extenders that results in the false detection of a WAN connectivity issue.
240 seconds	Interval after which Tunnel Manager declares the site-to-site secure tunnel destroyed due to missed heartbeats. In standalone mode, both the InterCloud Extender and InterCloud Switch will be rebooted. In HA mode, a switchover will occur and the failed InterCloud Link will reboot and come back as standby.

Management and Tunnel Interface Redundancy

The active InterCloud Extender sends a heartbeat request to the standby InterCloud Extender who then sends a reply. If the standby InterCloud Extender does not receive a heartbeat request for 30 seconds, it will explicitly send a request to the active InterCloud Extender. If no response is received, the standby InterCloud Extender sends a heartbeat request through the InterCloud Switch in its InterCloud Link who then forwards to its HA peer InterCloud Switch in the active InterCloud Link and finally to the intended active InterCloud Extender.

If a response is received, the InterCloud Extenders will print logs describing the detection of a connectivity issue between InterCloud Extenders. If no response is received, the standby InterCloud Extender will initiate a switchover. As a result, the standby InterCloud Link will move to the active state and the failed InterCloud Link will be rebooted and come up in the standby state.



Note

If the control connection between any of the gateways is broken, the redundant heartbeat mechanism will fail. If two redundant heartbeat requests fail across 5 seconds, the source InterCloud Extender will consider its HA peer InterCloud Extender failed.

Split Brain Resolution

When connectivity issues exist among InterCloud Extenders, Intercloud Switches, and the Cisco Prime Network Services Controller, it is possible that both InterCloud Extenders take the active role. This condition is called active-active or split-brain condition. When the communication is restored within the system, the InterCloud Extenders exchange information to decide which would have a lesser impact on the system, if rebooted.

A split-brain is not possible only due to InterCloud Extender connectivity loss because when the standby InterCloud Extender moves to active due to heartbeat failure, it sends a request to the Cisco Prime Network Services Controller to reboot the failed InterCloud Link. Once the failed InterCloud Link comes up it will move to the standby state.

A split-brain scenario is possible if there is a connectivity loss between the Cisco Prime Network Services Controller and the standby InterCloud Extender who is moving active. In this situation, the gateways in the failed InterCloud Link will not be rebooted, creating a active-active scenario.

If an active-active scenario occurs, the following parameters are considering during handshake resolution:

- Heartbeats Sent - The InterCloud Extender with a greater number of heartbeats sent within some threshold will remain active. If the difference in heartbeats sent is insignificant, the resolution will occur based on HA role.
- HA role - The InterCloud Extender with HA role Primary will remain active.

The InterCloud Link which moves to the standby state will be rebooted because many of the platform components do not support an active to standby state transition. The rebooted InterCloud Link will move to standby after it performs the handshake with the active InterCloud Extender.

Split Brain Resolution