

Building Highly Available Layer 3 Networks with Cisco NX-OS Software and Cisco Nexus 7000 Series Switches

White Paper

July, 2011

What You Will Learn	3
Background	3
Notes	3
High-Availability Best Practices for Layer 3 Networks and Recovery	3
System-Level High Availability and Best Practices	4
Software High Availability and Best Practices	4
Data-Plane High-Availability Best Practices	6
Protocol Fast Convergence Features	7
Protocol Fast Convergence: Importance of Timers	8
Fast Failure Detection	8
Fast Failure Reaction	9
General ISSU and Stateful Switchover Considerations	9
General vPC and Non-vPC Considerations	10
Bidirectional Forwarding Detection	10
Conclusion	13
Terminology	13
For More Information	14

What You Will Learn

In this document, you will learn about best practices for building highly available Layer 3 networks with Cisco® NX-OS Software and Cisco Nexus® 7000 Series Switches.

Background

As a result of data center trends such as consolidation and virtualization coupled with today's demanding applications, IT departments must help ensure a high level of network uptime. This challenge is complex and requires the right network design as well as following configuration best practices for different protocols. The Cisco NX-OS operating system and Cisco Nexus 7000 Series Switches provide many advanced features that help ensure continuous data center operation. This document describes those features and discusses best practices for deploying and configuring them, with the focus on Layer 3 protocol fast convergence.

Notes

For the purposes of this document only:

- The features discussed in this document are Cisco NX-OS features supported on the Cisco Nexus 7000 Series.
- Fast Failure Detection (FFD) and Fast Failure Reaction (FFR) may refer to separate features or to subfeatures of a particular protocol targeted for fast convergence.
- Layer 3 protocols refer to unicast routing protocols, multicast protocols as well as First-Hop Redundancy Protocols (FHRPs).
- FHRP refers to protocols such as Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP).

High-Availability Best Practices for Layer 3 Networks and Recovery

In information technology, high availability refers to a system or component that is continuously operational for a desirably long length of time. A highly available network is a network that can pass traffic for a long time period without any network-wide disruption of traffic. If a disruption does occur in a high-availability network, it is mitigated and isolated quickly so that end users do not see any effect. Cisco NX-OS and Cisco Nexus 7000 Series Switches can provide multiple levels of redundancy and recovery features so that the goals of high availability can be met.

The Cisco Nexus 7000 Series is a modular switching platform, with an emphasis on redundant critical components throughout all its subsystems. A highly modular and compartmentalized approach to systems design is applied to all facets of the platform, spanning the physical, environmental, power, and system software architectures. Additionally, a distinct functional separation between the control plane and forwarding data plane is emphasized in its design to allow continuous operation and no service disruption during planned or unplanned control-plane events or failures. This section presents features and best practices for system-level, software, and data-plane high availability and for protocol fast convergence.

System-Level High Availability and Best Practices

The Cisco Nexus 7000 Series Switches provide various system-level hardware redundancy features. These features include the following capabilities:

- Dual supervisor redundancy offers capabilities such as stateful supervisor switchover and provides the foundation for In-Service Software Upgrade (ISSU).
- Redundant fabric modules provide protection against individual fabric module failures.
- Redundant power supplies protect the system against either power supply failure or grid failure when configured appropriately.
- Network administrators have the ability to install multiple hardware I/O modules which allows them to build networks with path and I/O module diversity for PortChannel and equal-cost multipathing (ECMP) links.

Cisco highly recommends that hardware be configured with the maximum level of redundancy to provide the highest level of system availability.

Software High Availability and Best Practices

Cisco NX-OS provides a comprehensive approach to helping ensure control-plane and data-plane high availability. Cisco NX-OS uses a highly modularized architecture that compartmentalizes components for redundancy, fault isolation, and resource efficiency. Functional feature components operate as independent processes referred to as services. Cisco NX-OS services implement availability features by design into each service, as needed. Most system services are capable of performing stateful restarts, thereby allowing a given service experiencing a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network. Back-end management and orchestration of processes, services, and applications within a platform are handled by a set of high-level system-control services such as:

- System manager
- Persistent storage service
- Message and transaction service

These services also help ensure overall high availability. In a redundant configuration, such as when dual supervisor modules are in operation, mirrored services run on each supervisor module, with the configuration and operating state synchronized between them. One of these supervisors operates as the active supervisor, and the other operates in a hot-standby mode until activated in a switchover.

The system manager orchestrates overall system functions, service management, and system health monitoring. It is also responsible for maintaining overall high-availability states, enforcing high-availability policies, and managing system configuration. The system manager is responsible for launching, stopping, monitoring, and restarting services. The system manager is also used to initiate and manage the synchronization of service states and intersupervisor states for stateful switchover. The system manager will initiate a supervisor switchover if it determines that the current supervisor has undergone an unrecoverable failure or if critical core services are experiencing errors and cannot be restarted reliably.

In addition, as the overall control and monitoring process, the system manager is responsible for triggering the keep-alive indicator for the hardware-based watchdog timer on the supervisor. The lack of this periodic heartbeat from the system manager within the keep-alive timeout period of the watchdog timer indicates a nonresponsive system manager, which will trigger a hardware-based supervisor reset (single supervisor) or switchover (dual supervisors). The system manager's health is also monitored by a kernel-level module that receives periodic heartbeats sent by the system manager process. This monitoring allows the system to take corrective action in response to an unresponsive system manager that has exceeded the heartbeat timeout period.

As can be seen, Cisco NX-OS provides a multilevel high-availability architecture in which even high-availability components such as the system manager have protection in case they fail.

The persistent storage service (PSS) is the base infrastructure component responsible for storing and managing the operational run-time states of services. It is used by supported services to recover the runtime state in the event of a service restart. PSS serves as a database of state and run-time information. Services capable of using the PSS infrastructure can checkpoint their state information periodically. This feature, known as stateful restart, allows subsequent recovery to the last-known operating state preceding a failure. This state recovery capability is available to Cisco NX-OS services in both single and dual-supervisor configurations and helps enable a service to transparently return to operation without any impact on data-plane traffic, neighboring devices, or other internal services. For example, even in a single-supervisor configuration, the PSS enables the stateful restart of a service such as Spanning Tree Protocol without affecting the overall Spanning Tree Protocol topology or stability.

Routing protocols in Cisco NX-OS support protocol-specific graceful restart extensions. If a protocol cannot retrieve its run-time data stored in PSS for some reason or cannot use PSS, it can restart from a clean state by using graceful restart extensions. These extensions to base protocols provide nonstop forwarding (NSF) and least-obtrusive routing control-plane recovery. When a protocol that supports graceful restart starts from a clean state, it signals the neighbor to continue forwarding traffic while items pertaining to the run-time state such as neighbor relationship and Link-State Advertisement (LSA) databases are being built. The Cisco NX-OS routing protocol extensions for graceful restart follow the standards outlined in RFCs 3623, 5187, 4724 and 3847 for Open Shortest Path First protocol versions 2 and 3 (OSPFv2 and v3), Border Gateway Protocol Version 4 (BGPv4), and Intermediate System-to-Intermediate System (IS-IS) protocol, respectively. Note that although the Cisco NX-OS standards-based implementations of graceful restart extensions are compatible with Cisco IOS® Software releases that support the same IETF standards, they are not compatible with older Cisco IOS Software releases that only support the prestandard implementations of Cisco NSF.

In addition to the standards-based implementations for the protocols mentioned here, graceful restart extensions have been developed for use with Enhanced Interior Gateway Routing Protocol (EIGRP). The extensions to EIGRP used in Cisco NX-OS are compatible with those used for EIGRP NSF on other Cisco platforms.

Here are some best practices associated with Cisco NX-OS high availability as they pertain to Layer 3 networks:

- As mentioned in the discussion of hardware redundancy, you should maintain two supervisors in each system. Besides hardware redundancy mechanisms, you should use software redundancy mechanisms to help ensure that a particular process failure inside a single Cisco Nexus 7000 Series chassis does not affect the entire network.
- Use of routing protocols such as OSPF and IS-IS, which support stateful process restart, is recommended for better high availability. Stateful restart for other protocols, such as EIGRP and BGP, is planned.

- Regardless of whether the routing protocol can support stateful restart, make sure that graceful restart extensions always are enabled. They are enabled by default and should never be disabled. For example, for OSPF, use this command to verify that graceful restart extensions are enabled:

```
switch# show run ospf all
feature ospf
router ospf 1
...
    graceful-restart
    graceful-restart grace-period 60
...
```

Compatible graceful restart extensions need to be enabled on both peers to work properly.

- Taking Cisco NX-OS modularity into account, do not enable any feature in Cisco NX-OS unless you plan to use it. Processes associated with features that are not enabled will not be running and hence will not be taking up memory space.

Data-Plane High-Availability Best Practices

Cisco NX-OS provides two main features to achieve data-plane high availability: PortChannels and ECMP. The PortChannel feature in Cisco NX-OS allows users to aggregate up to 8 or 16 physical links (depending on the I/O module hardware and corresponding Cisco NX-OS Software release) for active-active redundancy and increased bandwidth. As a result of this aggregation, a single logical link is presented to Layer 2 or Layer 3 protocols. Traffic is load-shared across the links in a PortChannel using one of the available load-balancing schemes. For Layer 3 PortChannels, links can run only between a pair of devices; hence, only link-level redundancy can be achieved. For Layer 2 PortChannels, virtual PortChannels (vPCs) can be used to provide multichassis EtherChannel (MCEC) functions and, ultimately, chassis-level redundancy. Use of PortChannels helps ensure that failure of a single link does not cause networkwide effects. The failure instead will result in only a small traffic loss followed by rehashing of traffic to the remaining operational links. This process will continue until one of the following conditions occurs:

- The last operational link in the PortChannel fails.
- The number of operational links in the PortChannel falls below the [min-links](#) value specified by the administrator for Link Aggregation Control Protocol (LACP) for the PortChannel.
- The routing metric changes for the path going out from the Layer 3 PortChannel, making another path preferable. Metric changes can occur after any single link failure.

After one of these conditions occurs, traffic will need to switch to an alternative path if such a path is available.

ECMP can be used by Layer 3 protocols to provide link- and chassis-level redundancy as well as increased bandwidth. Note that ECMP is the only means of chassis-level redundancy for Layer 3 in Cisco NX-OS. As the name implies and contrary to PortChannel technology, ECMP allows Layer 3 protocols to see multiple paths to the destination. Losing one of the paths typically results in just a small packet loss, but it may also cause a routing event to affect more than just a pair of network devices. Each of the paths can be a PortChannel or a physical link. This is true as long as the path costs are equal. The number of paths available for traffic forwarding depends on the maximum number supported by Cisco NX-OS (currently, 16) as well as the number of paths that the routing protocol actually installs in the routing information base (RIB). For some protocols, the default number of installable paths is not the maximum number. Hence, the configuration may need to be modified to be able to install a larger number of paths, as, for example, in the case of OSPF:

```

switch(config)# router ospf 1
switch(config-router)# maximum-paths ?
<1-16> Maximum paths per destination
      *Default value is 8


```

On the basis of the information discussed in this section, following are best practices and considerations for achieving data-plane redundancy for Layer 3 networks:

- For link-level redundancy, use Layer 3 PortChannels or ECMP with at least two links available for a particular traffic destination.
- For chassis-level redundancy, use ECMP consisting of multiple links or PortChannels. Network architects must remember that, depending on routing protocol bandwidth calculations, a failure of a single link in a PortChannel can result in path failure. This behavior can be modified according to the amount of bandwidth that should be protected before the path is considered failed.
- For the two preceding best practices, always spread links in a PortChannel or ECMP links across at least two different I/O modules to use the high availability offered by a modular chassis platform. Use multiple modules for link-level redundancy first and then for chassis-level redundancy if additional modules are available.

Figure 1 shows one possible highly available design that takes into account the three points mentioned here.

Figure 1. Cisco Nexus 7000 Series Data-Plane Redundancy



Protocol Fast Convergence Features

In addition to system- and service-level high-availability capabilities, Cisco NX-OS on the Cisco Nexus 7000 Series provides several features to help protocols detect network failures and respond to and recover from those failures quickly, thus reducing the number of lost packets. Reducing the number of lost packets can help decrease network convergence time from seconds to milliseconds, which can make a difference in the overall effects seen by applications and application users. These features are discussed in more detail in the following sections.

While Cisco recommends that all high-availability and protocol fast convergence best practices be followed, you can choose to follow a subset based on your design and application tolerance requirements. For more information about Cisco NX-OS and Cisco Nexus 7000 Series high-availability features, see the “For More Information” section later in this document.

Protocol Fast Convergence: Importance of Timers

Fast convergence is a large and important topic in any modern network design. Typically, fast convergence refers to a collection of features as well as subfeatures of particular protocols used for faster-than-default network convergence in response to a failure. Fast convergence features can be categorized into two general types:

- Fast failure detection (FFD): Features that affect the speed with which a neighbor failure event can be detected
- Fast failure reaction (FFR): Features that affect the speed with which protocols notify the rest of the network that a change has occurred and the speed with which protocols running on network devices recalculate the new state based on an event that was received

Fast Failure Detection

Most modern Layer 3 protocols have a concept of neighbor relationships and exchange keepalives, heartbeats, or hellos, to maintain state. All Layer 3 protocols supported in Cisco NX-OS with exception of Routing Information Protocol Version 2 (RIPv2) use this mechanism. Hellos are used to establish initial communication with a neighbor and maintain a relationship. Typically, each protocol has a certain interval that neighbors use to exchange hellos. If no hellos are received within a period specified by a particular timeout value, referred to as dead time or hold time by some protocols, the neighbor is considered to be dead, and any routes whose next hop was pointing to the dead neighbor have to be recalculated. It is important that the neighbor down detection occur as quickly as possible and modifying hello and dead intervals may be a way to achieve that. For example, default hello and dead timers for OSPF are 10 seconds and 40 seconds, respectively. Waiting for 40 seconds to realize that a neighbor is down would not be an acceptable behavior for many customers. Hence, modifying the default hello and dead timeout values may be the way to achieve FFD, and so these modified timers can be referred to as FFD timers.

As with all features and settings, you should apply them only if they are needed. In deciding whether applying FFD timers will make a difference in neighbor-down detection time, consider a few general scenarios:

- Data-path failures: If Layer 3 connectivity between two peers is point to point and a physical link fails completely, applying FFD timers typically will not improve the convergence time. Convergence time will not be improved because on the Cisco Nexus 7000 Series and many other products, link-down notification to protocols will always be faster than default dead-timer expiration. However, if Layer 3 connectivity between two peers is not point to point, applying FFD timers will make a difference. Typically, this scenario occurs when a Layer 2 switch or coarse or dense wavelength-division multiplexing (CWDM or DWDM) multiplexer is in the data path of hellos.
- Control-path failures: As mentioned in prior sections, Cisco NX-OS provides a comprehensive high-availability infrastructure in which the system manager monitors processes and executes a process restart if it detects a failure. This infrastructure combined with the use of the other best practices discussed in this document should eliminate or reduce the impact on the network of most failure scenarios. In an extremely rare and unlikely case, the system manager may not detect any problems when a protocol running on the supervisor fails to send or receive hellos provided that the data path is operational. In such a case, reducing the hello interval and dead timeout value can reduce the time the system takes to detect a peer-down event.

This discussion applies to hello and dead-timer modifications for each protocol that needs to achieve FFD. When the number of protocols and interfaces grows large, you should consider other solutions to achieve FFD such as Bidirectional Forwarding Detection (BFD), discussed later in this document.

Fast Failure Reaction

As discussed here, the capability to detect a peer failure quickly is only a part of the solution to the problem of reducing network convergence time. The other important part of solving this problem is helping ensure that after a failure is detected, the network can recover quickly. FFR can be addressed by answering two questions. First, how quickly can the protocols running in network devices affected by a failure inform the rest of the network about what happened? For example, in the case of OSPF, you can modify the speed with which LSAs are generated. Second, how quickly can protocols running in all network devices recalculate their states based on a received notification? OSPF allows you to modify the speed with which the shortest-path first (SPF) tree is recalculated based on received LSAs. Depending on the particular Layer 3 protocol supported by Cisco NX-OS, features exist to address one or both of these questions. These features typically consist of timers, which can be referred to as FFR timers, and other settings. Details are provided in later sections.

For most networks, implementing some FFR features should be the first and most significant step in improving network convergence. Depending on network design and scalability, some FFR features may have a bigger impact on convergence than others, and some features may not have any impact unless they are configured in conjunction with other features. For example, if you quickly notify devices of a failure by generating LSAs in a fast manner but network devices are not recomputing the SPF tree quickly, the impact on convergence will be small at best. In fact, however, there are extremely few cases in which no single FFR feature or a combination of features can provide a meaningful improvement in convergence. Unless there is a guarantee that no failures will occur, administrators should consider implementing FFR features. However, there are still trade-offs to consider when implementing FFR timers and settings:

- With network changes affecting a very large number of routes, CPU use can increase greatly if FFR timers are set to very low values. That is why it is recommended that you test your configuration with the exact network conditions before implementing it in a production environment.
- You should maintain the same FFR timer and settings configuration in all devices in the network. If a device does not support a particular FFR feature, the networkwide configuration should be modified to the minimum supported level for all devices.

General ISSU and Stateful Switchover Considerations

Configuring FFD and FFR features has definite implications for stateful switchover (SSO) and ISSU. While disruptive supervisor switchovers and software upgrades can be achieved, Cisco does not support nondisruptive Cisco NX-OS Software upgrades and supervisor switchovers with modified FFD timers. In the case of ISSU, the system budgets a certain amount of time for processes to become active after switchover to a supervisor with a new code. Until processes are active, no hellos are sent and no run-time state is calculated.

If the dead timeout value on a switch expires before its peer's protocol process becomes active, a neighbor-down event may be registered and the network affected for no reason. This scenario is referred to as a false-positive event.

Although Cisco supports SSO and ISSU with FFR features, the factors mentioned in the FFR section may become important in cases in which a networkwide event occurs during ISSU. If the CPU is busy bringing up processes

after supervisor switchover, adding load by recalculating the SPF tree more quickly will certainly have an impact on the CPU load.

General vPC and Non-vPC Considerations

Implementation of the Cisco NX-OS vPC feature in a network triggers a few additional considerations for FFD timers for FHRPs. If vPC is configured at the boundary of Layer 2 and Layer 3 networks, FHRP behavior is modified to an active-active model with both nodes forwarding locally received traffic. Therefore, there is no need to modify hello interval and dead timeout value particularly for HSRP and VRRP since all nodes will already be forwarding traffic.

There are no known unique considerations regarding configuration of FFR and FFD features for routing protocols in a vPC scenario. All the other best practices mentioned previously can be used.

Bidirectional Forwarding Detection

As described earlier, FFD is a very important component of Layer 3 network convergence. However, modifying FFD timers for each protocol and for each interface on which that protocol runs can be problematic. Some important drawbacks of this approach are:

- Increased CPU load because the supervisor CPU must generate hellos more often for each protocol and for each interface on which that protocol runs. For example, if the CPU use goes up to 100 percent, a switch may miss generation or processing of hellos. Hence, more false positives will be seen in the network due to shorter hello and dead intervals.
- Operational challenges because the administrator must configure modified timers for many protocols on multiple devices and keep track of FFD timer configuration best practices for each protocol.
- Wasted link bandwidth because multiple sets of hellos are sent on the wire at a short interval by different protocols.

BFD is a protocol that tries to solve these problems. It implements a lightweight hello communication mechanism and builds a neighbor relationship state. This common mechanism, standardized in RFC 5580, can be used by multiple Layer 3 protocols for FFD instead of using per-protocol hellos for this purpose. Typically, Layer 3 protocols include a lot of additional information in hellos that is important for neighbor relationship setup but is not important for FFD. BFD hellos are much simpler messages that take less CPU power to generate and use less bandwidth. These hellos also are sent at a uniform rate, giving all protocols that use BFD the same predictable failure detection time. BFD is specified to run between two Layer 3 peers on top of any transport medium, including IPv4, IPv6, and Multiprotocol Label Switching (MPLS). The protocol supports active and passive roles and three operating modes: asynchronous, demand, and echo. BFDv0 and v1 are specified by IETF, with most modern implementations using BFDv1. Typically, BFD implementations run in a point-to-point fashion between two Layer 3 peers; however, some also implement [RFC 5883](#) to support BFD in a multi-hop environment.

General advantages of BFD include:

- Reduced CPU load because fewer per-protocol hellos need to be generated; each protocol sends initial hellos for neighbor discovery and synchronization, and after neighbors have been discovered, BFD hellos are used for keepalive and failure detection purposes.
- Increased link bandwidth availability because only BFD sends hellos at a shorter interval while other protocols continue sending hellos at the default interval.

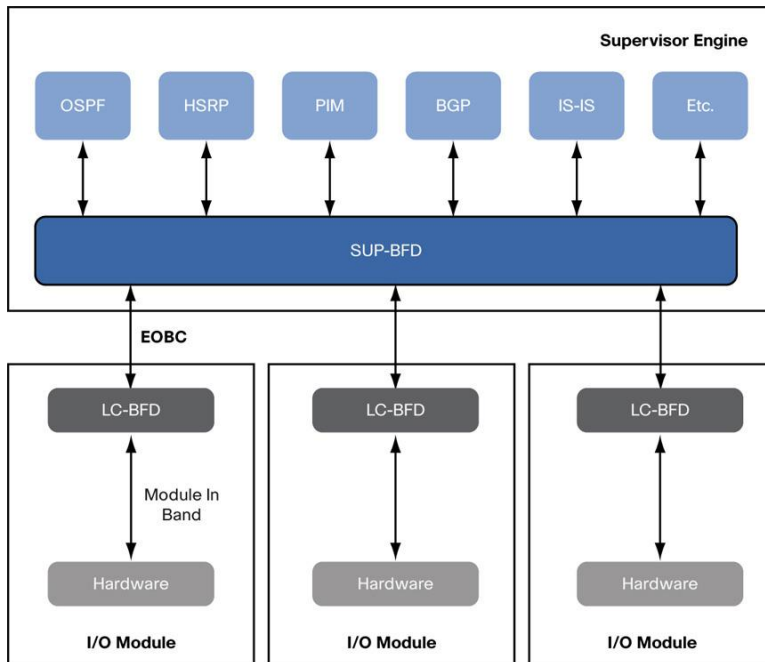
-
- Failure detection in less than a second can be provided; this BFD capability is especially important for certain protocol implementations that do not support sending hellos at intervals of less than a second.
 - Reduced operational complexity because hello and dead timers do not have to be configured individually.

The Cisco Nexus 7000 Series supports BFDv1 starting with Cisco NX-OS Software Release 5.0.2a. Only asynchronous mode with or without echo is supported in a single-hop environment. Several protocols are supported to take advantage of BFD in Cisco NX-OS, including OSPFv2, EIGRP, HSRP, Protocol-Independent Multicast (PIM) and BGP. BFD can also run on many types of interfaces, including physical interfaces, Layer 2 and Layer 3 PortChannels, and subinterfaces. Be sure to check the Cisco NX-OS configuration guides for Cisco Nexus 7000 Series and release notes for information about the latest supported configurations.

Cisco NX-OS has a sophisticated distributed implementation of BFD, shown in Figure 2. SUP-BFD is a BFD process that runs in the supervisor and is responsible for communication with protocols such as OSPF and HSRP that can request the use of BFD services. Those protocols become clients of BFD. LC-BFD is a process that is hosted by CPUs of individual line cards and is responsible for generating and receiving actual hellos. A particular client protocol such as OSPF may establish a neighbor relationship with a peer. After that relationship is established, OSPF registers with BFD, and SUP-BFD sends a message across the Ethernet out-of-band channel (EOBC) to a line card. That message instructs LC-BFD to establish a BFD session with a peer. Peers use a three-way handshake mechanism to establish a session. Through that mechanism, the detect multiplier, session descriptor, and receive (Rx) and transmit (Tx) intervals for BFD hellos are negotiated.

After a session is established, LC-BFD notifies SUP-BFD, which in turn notifies the client protocol that registered with BFD. From this point, BFD hellos can be used for forwarding-path failure detection. BFD hellos are sent using a user-specified interval with a 50-millisecond minimum and default. If a number of hello packets are missed (the default is 3), the BFD session is declared dead. When the BFD session goes down, LC-BFD notifies SUP-BFD, which in turn notifies the client protocol. After that, each protocol is free to take an independent action based on information provided by BFD.

Figure 2. Cisco NX-OS BFD Software Architecture on Cisco Nexus 7000 Series



The Cisco NX-OS BFD architecture and implementation has several advantages:

- Offloading the most intensive and repetitive tasks of hello generation to I/O module CPUs further reduces the load on the supervisor CPU.
- Use of echo mode, enabled by default, allows the switch to reduce the use of I/O module CPUs because the recipient will be able to use the hardware replication engine to return the same hello packet to the sender.
- For many subinterfaces belonging to a single physical interface, a single master BFD session can be run for one subinterface at a hello interval specified by the user, and slave sessions can be run for other subinterfaces at a higher hello interval. This feature further reduces the load on I/O module CPUs.
- The BFD hello and dead intervals do not have the same SSO and ISSU considerations as per-protocol FFD timers do. Because a process that generates BFD hellos runs on an I/O module, it is not affected by supervisor switchovers and software upgrades. The SUP-BFD process undergoes a stateful restart, which will not trigger any false positives for neighbor-down events.
- Per-link BFD sessions can be run in a Layer 3 PortChannel for greater resiliency.
- BFD is virtual routing and forwarding (VRF) and virtual device context (VDC) aware, allowing it to run on interfaces allocated to specific VRF instances and specific VDCs.

The following example shows BFD configuration for OSPF, assuming that OSPF is already configured:

```
switch(config)# feature bfd
switch(config)# bfd interval <50-999> min_rx <50-999> multiplier <1-50>
switch(config)# router ospf 1
switch(config-router)# bfd
```

As of Cisco NX-OS 5.0.2(a), BFD has been tested to support up to 200 sessions at a 50-millisecond interval times 3 (multiplier) per I/O module, allowing FFD in as little as 150 milliseconds. These timers are defaults for BFD, and you should not change them unless you are concerned about link quality and too many false positives.

Considering all the general benefits of the BFD protocol as well as the unique advantages of BFD implementation on Cisco NX-OS and the Cisco Nexus 7000 Series, Cisco strongly recommends the use of BFD instead of per-protocol FFD timers.

Conclusion

When designing a new Layer 3 network, it is important to understand all system, software, and data-plane high-availability and protocol fast convergence features that are offered by your platform of choice. This document presented an overview of the advantages and best practices for building highly available Layer 3 networks with Cisco NX-OS Software on Cisco Nexus 7000 Series Switches. By strategically taking advantage of product capabilities and following the best practices discussed in this document, customers can build scalable and self-healing networks that meet their current and future business needs.

Terminology

Table 1 defines acronyms used in this document.

Table 1. Acronyms

Acronym	Definition
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
DWDM	Dense Wave-Division Multiplexing
ECMP	Equal-cost multipathing
EOBC	Ethernet out-of-band channel
FFD	Fast failure detection
FFR	Fast failure reaction
FHRP	First Hop Redundancy Protocol
GLBP	Gateway Load Balancing Protocol
GR	Graceful Restart
HA	High Availability
HSRP	Hot Standby Router Protocol
IETF	Internet Engineering Task Force
ISSU	In-Service Software Upgrade
LSA	Link-state advertisement
NSF	Non-stop Forwarding
OSPF	Open Shortest Path First
RIB	Routing information base
RFC	Request for comments
SSO	Stateful switchover
STP	Spanning Tree Protocol
VDC	Virtual device context
VRRP	Virtual Router Redundancy Protocol

VPC	Virtual PortChannel
-----	---------------------

For More Information

- [Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x](#)
- [Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x](#)
- [Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 5.x](#)
- [Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x](#)
- [Cisco Nexus 7000 Switches: Continuous Operations and High Availability](#)
- [Cisco IOS Software Releases 12.2 SX: Bidirectional Forwarding Detection](#)
- [Cisco IOS IP Routing: BFD Configuration Guide, Release 15.1](#)
- [Cisco IOS IP Routing: BFD Configuration Guide, Release 12.2SR](#)
- [Cisco IOS IP Routing: BFD Configuration Guide, Release 12.2SX](#)
- [Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Configuration Guide: BFD](#)
- [Cisco IOS XR Interface and Hardware Component Configuration Guide for the Cisco CRS Router, Release 4.0](#)
- [IETF BFD Working Group](#)



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Printed in USA

C11-673862-00 07/11