Configuring High Availability

This chapter describes Stateful Switchover, Cisco Nonstop Forwarding, and In-Service Software Upgrade. This chapter also describes the configuration procedures.

This chapter includes the following topics:

Stateful Switchover

Stateful switchover (SSO) ensures state synchronization and non-disruptive switchover from an active to a standby fabric card, thereby providing an increase in both system and network availability. In SSO, the standby fabric card is fully initialized and is ready to assume control from the active fabric card when the switchover occurs.

SSO is particularly useful at the network edge. Traditionally, core routers protect against network faults using router redundancy and mesh connections that allow traffic to bypass failed network elements. SSO provides protection for network edge devices with dual route processors (RPs) that represent a single point of failure in the network design, and where an outage might result in loss of service for customers.

Because SSO maintains stateful protocol and application information, the user session information is maintained during a switchover. Line cards continue to forward network traffic with no loss of sessions, providing improved network availability.

SSO provides faster switchover by fully initializing and configuring the standby fabric card, and by synchronizing state information, which reduces the time required for routing protocols to converge. The CPT 50 panel that is connected directly to the active fabric card will be disconnected due to SSO.

SSO takes advantage of fabric card redundancy to increase network availability. SSO establishes one of the fabric cards as the active card and the other fabric card is established as the standby card, and then synchronizes critical state information between them. Following an initial synchronization between the two cards, SSO dynamically maintains state information between them.

SSO enables the system to keep all the sessions of HA-aware protocols up during a switchover. This ensures that the services are not affected during switchovers and upgrades.

Note Carrier Packet Transport (CPT) supports only the SSO redundancy mode.

Cisco Nonstop Forwarding (NSF) works with SSO to minimize the network downtime following a switchover. The main objective of NSF is to continue forwarding IP packets following a fabric card switchover.

Prerequisites for SSO

Two fabric cards must be installed on the CPT 200 and CPT 600 shelf.

Restrictions for SSO

General Restrictions

For NSF support, neighboring routers must run NSF-enabled images, though SSO need not be configured on the neighboring device.

Switchover Process Restrictions

If the active fabric card fails before the standby fabric card is ready to switchover, both the fabric cards are reset.

Platform Restrictions or Considerations

- CPT supports only the SSO redundancy mode.
- The active fabric card reload operation causes a switchover.

Card Crash During SSO

When the fabric card or line card crashes during SSO, the crash handler shuts down the laser signal on the front ports of the card. The far end detects loss of signal on this path and switches to the protection path if available.

SSO Synchronization

Fabric Card Synchronization

Both the fabric cards must be running the same configuration so that the standby fabric card is always ready to assume control if the active fabric card fails.

To achieve the benefits of SSO, synchronize the configuration information from the active fabric card to the standby fabric card at startup and whenever changes to the active fabric card configuration occur. This synchronization occurs in two separate phases:

- When the standby fabric card is booting, the configuration information is synchronized in bulk from the active fabric card to the standby fabric card.
- When configuration or state changes occur, an incremental synchronization is conducted from the active fabric card to the standby fabric card.

Note The CPT 50 panels connected to the active fabric card and the active fabric card are reset during the switchover.

Bulk Synchronization During Initialization

When a system with SSO is initialized, the active fabric card performs a chassis discovery (discovery of the number and type of line cards and fabric cards in the system) and parses the startup configuration file.

The active fabric card then synchronizes this data to the standby fabric card and instructs the standby fabric card to complete its initialization. This method ensures that both fabric cards contain the same configuration information.

Even though the standby fabric card is fully initialized, it interacts only with the active fabric card to receive incremental changes to the configuration files as they occur.

Incremental Synchronization

After both the fabric cards are fully initialized, any further changes to the running configuration or active states are synchronized to the standby fabric card as they occur. Active fabric card states are updated as a result of processing protocol information, external events (such as the interface coming up or down), user configuration commands (using CLI commands or CTC), or other internal events.

Line Card State Synchronization

Changes to the line card states are synchronized to the standby fabric card. Line card state information is initially obtained during bulk synchronization of the standby fabric card. Following bulk synchronization, line card events received at the active fabric card are synchronized to the standby fabric card.

The following line card states apply to line cards during switchover:

- Line cards are not reset during switchover.
- Line cards are not reconfigured during switchover.
- Line cards do not generate alarms during switchover.
- Subscriber sessions are not lost during switchover.

Note During the complete power cycle of the fully loaded CPT 200 or CPT 600 chassis with large scale service configurations, sometimes the state of the cards will show as Empty Slot. The line cards boot properly once the database is completely loaded.

Counters and Statistics Synchronization

The various counters and statistics maintained in the active fabric card are reset during a switchover and are not synchronized with the standby fabric card.

Redundancy Modes

CPT supports only the SSO redundancy mode.

SSO Redundancy Mode

In SSO redundancy mode, the standby fabric card is fully initialized. The active fabric card dynamically synchronizes the startup and running configuration changes to the standby fabric card, which means that the standby fabric card need not be reloaded and reinitialized. SSO supports synchronization of line card, protocol, and application state information between the fabric cards.

Switchover Operation

During switchover, the system control and routing protocol execution are transferred from the active fabric card to the standby fabric card.

Switchover Conditions

The following conditions cause a switchover from the active fabric card to the standby fabric card:

- The active fabric card fails.
- Online removal of the active fabric card automatically forces a stateful switchover to the standby fabric card.
- Forced switchover from the active fabric card to the standby fabric card through CTC or CLI.
- The **reload** command causes a switchover from the active fabric card to the standby fabric card.
- Keepalive or Heartbeat failure causes a switchover from the active fabric card to the standby fabric card.

  The fabric card polls the peer cards at regular intervals to check whether the cards are connected to the system. If the fabric card does not receive a message within a specific duration from its peer card, the fabric card assumes that the peer card is down. The default duration for the fabric card to receive the Keepalive indication message is 9000 milliseconds.

Switchover Exceptions

The following conditions will not cause a switchover from the active fabric card to the standby fabric card:

- The standby fabric card reset doest not cause a switchover from the active fabric card to the standby fabric card.
- The Dual TNC card reset does not cause a switchover from the active fabric card to the standby fabric card.

NTP-J17 Manage SSO

Purpose	This procedure manages SSO configuration.
Tools/Equipment	CPT 600 / CPT 200
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote

Security Level	Provisioning or higher

Procedure

Perform any of the following procedures as needed:

- DLP-J46 Perform a Manual Switchover Using CTC, on page 5
- DLP-J47 Verify SSO Configuration Using Cisco IOS Commands, on page 6
- DLP-J48 Verify SSO Configuration Using CTC, on page 7
- DLP-J49 Troubleshoot SSO Using Cisco IOS Commands, on page 11

Stop. You have completed this procedure.

DLP-J46 Perform a Manual Switchover Using CTC

Purpose	This procedure performs a manual switchover from the active fabric card to the standby fabric card using CTC.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

The fabric card can be in one of the following states during the switchover.

- Active
- Failed
- Loading
- Standby

The active fabric card moves to one of the above states after the switchover.

Procedure

Step 1 Complete the NTP-J22 Log into CTC procedure at a node where you want to switchover from the active fabric card to the standby fabric card.

Step 2 From the View menu, choose **Go to Home View**.

Step 3 Right-click the fabric or line card and choose **Open Packet Transport System View**. The Packet Transport System View dialog box appears.

Step 4 Click the **Maintenance** tab.

Step 5 From the left pane, click **High-Availability**.

Step 6 In the Fabric Cards area, click **Switch** to perform a manual switchover from the active fabric card to the standby fabric card.

Step 7 Click **YES**.

Step 8 Return to your originating procedure (NTP).

DLP-J47 Verify SSO Configuration Using Cisco IOS Commands

Purpose	This procedure allows you to verify that SSO is configured using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**show redundancy** [**clients** | **config-sync** | **counters** | **domain** | **history** | **idb-sync-history** | **interlink** | **states** | **switchover** | **trace**] **Example:** Router# show redundancy	Verifies that SSO is configured on the networking device.

	Command or Action	Purpose
Step 3	Return to your originating procedure (NTP).	—

DLP-J48 Verify SSO Configuration Using CTC

Purpose	This procedure verifies SSO configuration using CTC.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

Procedure

Step 1 Complete the NTP-J22 Log into CTC procedure at a node where you want to verify the SSO configuration.

Step 2 From the View menu, choose **Go to Home View**.

Step 3 Right-click the fabric or line card and choose **Open Packet Transport System View**. The Packet Transport System View dialog box appears.

Step 4 Click the **Maintenance** tab.

Step 5 From the left pane, click **IOS**.

Step 6 Click **Open IOS Connection**. The IOS Login dialog box appears.

Step 7 Enter the user name and password.

Step 8 Enter show redundancy command at the prompt.

Step 9 Press **Enter**. The output is displayed.
The redundant system information, current processor, and peer processor information appears in the output area. The configured redundancy mode is SSO if an entry in the output reads as `Configured Redundancy Mode = SSO`.

Step 10 Return to your originating procedure (NTP).

Examples of SSO Configuration

Verify that SSO Is Configured

The following examples verify that SSO is configured on the device.

```
Router# show redundancy
```

```
Redundant System Information :
------------------------------
       Available system uptime = 18 hours, 44 minutes
       Switchovers system experienced = 1
       Standby failures = 0
       Last switchover reason = active unit failed

       Hardware Mode = Duplex
       Configured Redundancy Mode = SSO
       Operating Redundancy Mode = SSO
       Maintenance Mode = Disabled
       Communications = Up

Current Processor Information :
------------------------------
        Active Location = slot 5
        Current Software state = ACTIVE
        Uptime in current state = 10 minutes
        Image Version = Cisco IOS Software, ONS NGXP Software
(NGXP-ADVIPSERVICES-M), Experimental Version
        15.1(20110216:101154) [ios_ngxp_dev-georgeti-ios_ngxp_dev.pkg
100]
        Copyright (c) 1986-2011 by Cisco Systems, Inc.
        Compiled Wed 16-Feb-11 16:59 by georgeti
        Configuration register = 0x101

Peer Processor Information :
----------------------------
        Standby Location = slot 4
        Current Software state = STANDBY HOT
        Uptime in current state = 8 minutes
        Image Version = Cisco IOS Software, ONS NGXP Software
(NGXP-ADVIPSERVICES-M), Experimental Version
        15.1(20110215:170703) [ios_ngxp_dev-sathk-ngxp_Feb16th 109]
        Copyright (c) 1986-2011 by Cisco Systems, Inc.
        Compiled Wed 16-Feb-11 15:12 by sathk
        Configuration register = 0x101 (will be 0x8001 at next reload)
```

```
Router# show redundancy states
```

```
     my state = 13  -ACTIVE
     peer state = 8 -STANDBY HOT
     Mode = Duplex
     Unit ID = 4

    Redundancy Mode (Operational) = SSO
    Redundancy Mode (Configured)  = SSO
```

```
    Redundancy State               = SSO
    Manual Swact = enabled
    Communications = Up

    client count = 47
    client_notification_TMR = 30000 milliseconds
    keep_alive TMR = 9000 milliseconds
    keep_alive count = 0
    keep_alive threshold = 10
    RF debug mask = 0x0
```

```
Router# show redundancy history
```

```
00:00:12 client added: Redundancy Mode RF(29) seq=60
00:00:12 client added: IfIndex(139) seq=61
00:00:12 client added: CHKPT RF(25) seq=68
00:00:12 client added: NGXP Platform RF(4500) seq=76
00:00:12 client added: NGXP CardIntf Mgr RF(4505) seq=77
00:00:12 client added: Event Manager(77) seq=84
00:00:12 client added: Network RF Client(22) seq=109
00:00:12 client added: XDR RRP RF Client(71) seq=135
00:00:12 client added: CEF RRP RF Client(24) seq=136
00:00:12 client added: RFS RF(520) seq=157
00:00:12 client added: Config Sync RF client(5) seq=159
```

```
Router# show redundancy switchover history
```

```
Index  Previous  Current  Switchover         Switchover
       active    active   reason             time
-----  --------  -------  ----------------   --------------------
   1       4        5     active unit failed 10:58:11 PDT Wed Jun 7 2000
```

Verify SSO Protocols and Applications

Enter the **show redundancy** command with the **client** keyword to display the list of applications and protocols that have registered as SSO protocols or applications. You can also verify the list of supported line protocols.

```
Router# show redundancy clients
```

```
 clientID = 29       group_id = 1    clientSeq = 60       Redundancy Mode
 RF
 clientID = 139      group_id = 1    clientSeq = 61       IfIndex
 clientID = 25       group_id = 1    clientSeq = 68       CHKPT RF
 clientID = 4500     group_id = 1    clientSeq = 76       NGXP Platform
RF
 clientID = 4505     group_id = 1    clientSeq = 77       NGXP CardIntf
Mgr RF
 clientID = 4504     group_id = 1    clientSeq = 78       NGXP PB Manager

 clientID = 4501    group_id = 1    clientSeq = 79      NGXP HAL Resource
 Manager
 clientID = 4502     group_id = 1    clientSeq = 80       NGXP NMS Manager

 clientID = 4503     group_id = 1    clientSeq = 81       NGXP Mac Table
 Management
 clientID = 77       group_id = 1   clientSeq = 84       Event Manager
```

```
 clientID = 78       group_id = 1    clientSeq = 106      TSPTUN HA
 clientID = 22       group_id = 1    clientSeq = 109      Network RF
Client
 clientID = 75       group_id = 1    clientSeq = 126      Tableid HA
 clientID = 71       group_id = 1    clientSeq = 135      XDR RRP RF
Client
 clientID = 24       group_id = 1    clientSeq = 136      CEF RRP RF
Client
 clientID = 146      group_id = 1    clientSeq = 138      BFD RF Client
 clientID = 402      group_id = 1    clientSeq = 156      TPM RF client
 clientID = 520      group_id = 1    clientSeq = 157      RFS RF
 clientID = 5        group_id = 1    clientSeq = 159      Config Sync RF
 client
 clientID = 49       group_id = 1    clientSeq = 181      HDLC
 clientID = 72       group_id = 1    clientSeq = 182      LSD HA Proc
 clientID = 113      group_id = 1    clientSeq = 183      MFI STATIC HA
Proc
 clientID = 290      group_id = 1    clientSeq = 184      MPLS TP HA
 clientID = 204      group_id = 1    clientSeq = 188      ETHER INFRA RF
 clientID = 226      group_id = 1    clientSeq = 196      LACP
 clientID = 20       group_id = 1    clientSeq = 203      IPROUTING NSF
RF client
 clientID = 34       group_id = 1    clientSeq = 218      SNMP RF Client
 clientID = 35       group_id = 1    clientSeq = 228      History RF
Client
 clientID = 90      group_id = 1    clientSeq = 240      RSVP HA Services

 clientID = 54       group_id = 1    clientSeq = 256      SNMP HA RF
Client
 clientID = 73       group_id = 1    clientSeq = 257      LDP HA
 clientID = 76       group_id = 1    clientSeq = 258      IPRM
 clientID = 57       group_id = 1    clientSeq = 259      ARP
 clientID = 83       group_id = 1    clientSeq = 293      AC RF Client
 clientID = 82       group_id = 1    clientSeq = 294      CCM RF
 clientID = 84       group_id = 1    clientSeq = 296      AToM manager
 clientID = 85       group_id = 1    clientSeq = 298      SSM
 clientID = 280      group_id = 1    clientSeq = 299      ST PW OAM
 clientID = 212      group_id = 1    clientSeq = 309      REP Protocol
 clientID = 151      group_id = 1    clientSeq = 322      IP Tunnel RF
 clientID = 94       group_id = 1    clientSeq = 323      Config Verify
RF client
 clientID = 506      group_id = 1    clientSeq = 327      Igmp Snooping
 clientID = 3099     group_id = 1    clientSeq = 347      ISSU process
 clientID = 4005    group_id = 1    clientSeq = 350     ISSU Test Client

 clientID = 93       group_id = 1    clientSeq = 354      Network RF 2
Client
 clientID = 141     group_id = 1    clientSeq = 364      DATA DESCRIPTOR
 RF CLIENT
 clientID = 4020     group_id = 1    clientSeq = 393      IOS Config
ARCHIVE
 clientID = 4021     group_id = 1    clientSeq = 394      IOS Config
ROLLBACK
```

Possible SSO Problem Situations

This section describes the possible situations in which SSO troubleshooting may be needed.

- The standby fabric card was reset, but no error message was displayed. To display a log of SSO events and clues as to why a switchover or other event occurred, enter the **show redundancy history** command on the newly active fabric card:

  Router# **show redundancy history**

DLP-J49 Troubleshoot SSO Using Cisco IOS Commands

Purpose	This procedure allows you to troubleshoot the SSO feature using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

The following commands do not have to be entered in any specific order.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**crashdump-timeout** [*mm* | *hh:mm*] **Example:** router(config-red)# crashdump-timeout	Sets the longest time that the newly active fabric card waits before reloading the previously active fabric card.
Step 3	**show redundancy** [**clients** | **config-sync** | **counters** | **domain** | **history** | **idb-sync-history** | **interlink** | **states** | **switchover** | **trace**] **Example:** Router# show redundancy states	Displays the redundancy configuration mode of the fabric card. Also, displays information about the number of switchovers, system uptime, processor uptime, and redundancy state, and reasons for any switchovers.

	Command or Action	Purpose
Step 4	**show version** **Example:** Router# show version	Displays the image information for each fabric card.
Step 5	Return to your originating procedure (NTP).	—

SSO Aware Protocols and Applications

Protocols and applications that support SSO must be HA-aware. A feature or protocol is HA-aware if it maintains, either partially or completely, undisturbed operation through a fabric card switchover. State information for HA-aware protocols and applications is synchronized from the active fabric card to the standby fabric card to achieve stateful switchover for those protocols and applications.

The following protocols and applications are SSO aware:

- EVC, MVR, IGMP Snooping, MAC Address Table
- REP
- QoS
- LAG and LACP
- OSPF
- MPLS, LDP, RSVP-TE, L2VPN, BFD

Active-Active Data Path

Active-Active Data Path (AADP) is supported. AADP refers to the load sharing between the two fabric cards.

The redundant fabric cards run in an active-standby control model. However, both the fabric cards have ports that carry active traffic.

The active and standby fabric cards are used for load balancing the traffic that is received and destined for line cards. With dual fabric cards, the traffic from one line card to the other line card can be up to 40 Gbps. With a single fabric card, the traffic from one line card to the other line card cannot exceed 32 Gbps. However, the traffic between two fabric cards can be up to 32 Gbps for CPT 600 shelf, and up to 40 Gbps for CPT 200 shelf.

AADP provides the following benefits for the CPT System:

- Reduces the overall cost per 10 GE port.
- Increases the trunk capacity of the CPT system.
- Builds redundancy with rings or tunnels across the fabric cards.

Nonstop Forwarding

Cisco Express Forwarding

A key element of NSF is packet forwarding. In Cisco networking devices, packet forwarding is provided by Cisco Express Forwarding (CEF). Cisco Express Forwarding maintains the Forwarding Information Base (FIB), and uses the FIB information that was current at the time of the switchover to continue forwarding packets during a switchover. This feature eliminates downtime during the switchover.

During the normal NSF operation, Cisco Express Forwarding on the active fabric card synchronizes its FIB and adjacency databases with the FIB and adjacency databases on the standby fabric card. On switchover of the active fabric card, the standby fabric card initially has FIB and adjacency databases that are mirror images of those on the active fabric card. The packet forwarding continues after a switchover as soon as the interfaces and a data path are available.

Nonstop Forwarding

Nonstop Forwarding (NSF) works with the Stateful Switch Over (SSO) feature to minimize the amount of time a network is unavailable following a switchover. The main objective of the NSF is to continue forwarding IP packets after the switchover of the active fabric card.

When a networking device restarts, all the routing peers of that device usually detect that the device went down and then came back up. This down-to-up transition results in routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. NSF helps to suppress routing flaps and improves the network stability.

NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. With NSF, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards while the standby fabric card assumes control from the failed active fabric card during a switchover. The ability of line cards to remain up during the switchover and to be kept current with the FIB on the active fabric card is key to NSF operation. The CPT 50s connected to the fabric card will have an impact on the traffic.

Note CPT does not support forwarding IP packets in hardware and supports forwarding only in software.

The following figure illustrates how SSO is typically deployed in service provider networks. In this example, NSF with SSO is primarily at the access layer (edge) of the service provider network. A fault at this point can result in loss of service for enterprise customers requiring access to the service provider network.

Figure 1: NSF with SSO Network Deployment: Service Provider Networks

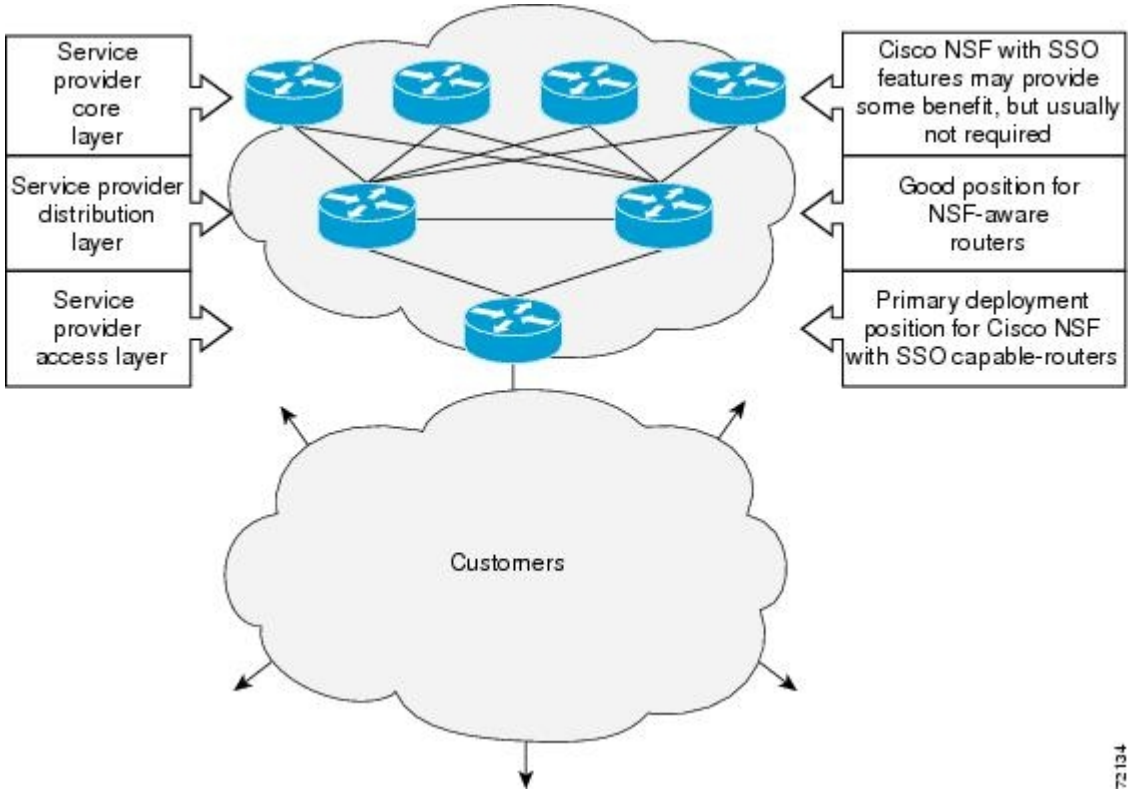

Additional levels of availability may be gained by deploying NSF with SSO at other points in the network where a single point of failure exists. The following figure illustrates an optional deployment strategy that applies NSF with SSO at the enterprise network access layer. In this example, each access point in the enterprise

network represents another single point of failure in the network design. In the event of a switchover or a planned software upgrade, enterprise customer sessions would continue uninterrupted throughout the network.

Figure 2: NSF with SSO Network Deployment: Enterprise Networks

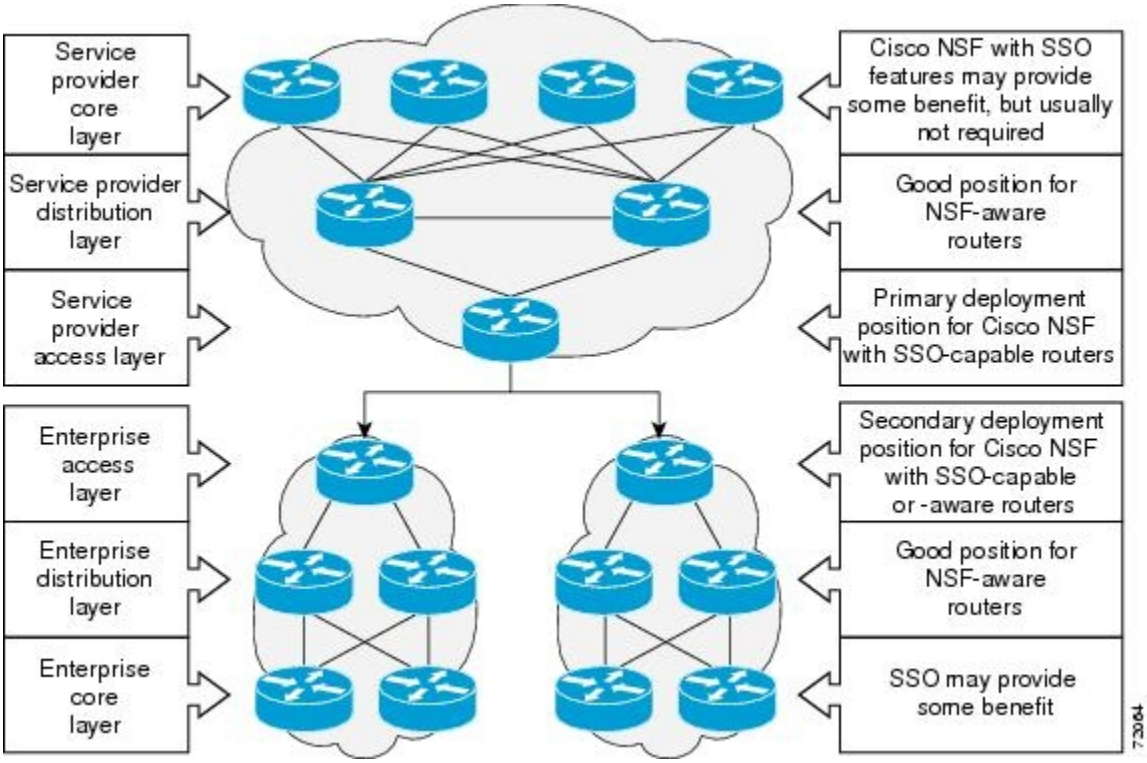

Prerequisite for NSF

- Ensure that Distributed Cisco Express Forwarding is operational.

Benefits of NSF

The NSF feature has the following benefits:

- Improved network availability—NSF continues to forward network traffic and application state information so that user session information is maintained after a switchover.
- Overall network stability—Network stability may be improved with the reduction in the number of route flaps that had been created when routers in the network failed and lost their routing tables.
- Prevents routing flaps—Because the NSF continues to forward network traffic in the event of a switchover, routing flaps are avoided.
- No loss of user sessions—User sessions established prior to the switchover are maintained.

NSF Routing

Each protocol depends on Cisco Express Forwarding to continue forwarding packets during switchover while the routing protocols rebuild the Routing Information Base (RIB) tables. After the routing protocols have

converged, Cisco Express Forwarding updates the FIB table and removes stale route entries. Cisco Express Forwarding updates the line cards with the new FIB information.

The routing protocols run only on the active fabric card, and they receive routing updates from the neighboring routers. After a switchover, the routing protocols request that the NSF-aware neighboring devices send the state information to help rebuild the routing tables.

Note During the NSF operation, the routing protocols depend on Cisco Express Forwarding to continue forwarding the packets while the routing protocols rebuild the routing information.

OSPF Operation

When an OSPF NSF-capable router performs a fabric card switchover, it must perform two tasks to resynchronize its link state database with its OSPF neighbors. Firstly, it must relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship. Secondly, it must reacquire the contents of the link state database for the network.

After the fabric card switchover, the NSF-capable router sends an OSPF NSF signal to the neighboring NSF-aware devices. The neighboring network devices recognize this signal as a clue that the neighbor relationship with this router must not be reset. As the NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

When neighbor relationships are reestablished, the NSF-capable router begins to resynchronize its database with all of its NSF-aware neighbors. At this point, the routing information is exchanged between the OSPF neighbors. After this exchange is complete, the NSF-capable device uses the routing information to remove stale routers and update the RIB and FIB with the new forwarding information. The OSPF protocols are then fully converged.

Note The OSPF NSF requires that all the neighboring network devices be NSF-aware. If a NSF-capable router discovers that it has non NSF-aware neighbors on a particular network segment, it will disable NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware routers will continue to provide NSF capabilities.

NSF Device Modes

Cisco had implemented the proprietary Cisco NSF. The Graceful OSPF Restart feature supports IETF NSF for OSPF processes in multivendor networks. The NSF device modes of operation common to the Cisco and IETF NSF implementations are as follows:

- Restarting mode—In this mode, the OSPF device is performing nonstop forwarding recovery because of the fabric card switchover.
- Helper mode—Also known as NSF–awareness. In this mode, the neighboring device is restarting and helping in the NSF recovery.

NTP-J18 Manage NSF

Purpose	This procedure manages NSF configuration.

Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

Procedure

Perform any of the following procedures as needed:

- DLP-J50 Verify Cisco Express Forwarding NSF Using Cisco IOS Commands, on page 17
- DLP-J51 Configure OSPF for NSF Using Cisco IOS Commands, on page 18
- DLP-J52 Verify OSPF for NSF Using Cisco IOS Commands, on page 19
- DLP-J53 Troubleshoot NSF Using Cisco IOS Commands, on page 21

Stop. You have completed this procedure.

DLP-J50 Verify Cisco Express Forwarding NSF Using Cisco IOS Commands

Purpose	This procedure verifies that Cisco Express Forwarding is NSF capable using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

The Cisco Express Forwarding NSF feature operates by default. Configuration of Cisco Express Forwarding NSF is not necessary.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.

	Command or Action	Purpose
Step 2	**show cef state** **Example:** Router# show cef state	Displays the state of Cisco Express Forwarding on a networking device.
Step 3	Return to your originating procedure (NTP).	—

DLP-J51 Configure OSPF for NSF Using Cisco IOS Commands

Purpose	This procedure configures OSPF for NSF using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** Router# configure terminal	Enters global configuration mode.
Step 3	**router ospf** *process-id* **Example:** Router(config)# router ospf 12	Enables an OSPF routing process, and places the router in router configuration mode.
Step 4	**nsf** {**cisco** | **ietf**} **Example:** Router(config-router)# nsf cisco	Enables Cisco NSF or IETF NSF support on the router.

	Command or Action	Purpose
Step 5	Return to your originating procedure (NTP).	—

DLP-J52 Verify OSPF for NSF Using Cisco IOS Commands

Purpose	This procedure verifies OSPF for NSF using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**show running-config** **Example:** Router# show running-config	Displays the contents of the current running configuration file.
Step 3	**show ip ospf nsf** **Example:** Router# show ip ospf nsf	Displays NSF information about OSPF routing processes.
Step 4	Return to your originating procedure (NTP).	—

Examples of the NSF Configuration

Verify that Cisco Express Forwarding Is NSF-capable

The following example shows how to verify that Cisco Express Forwarding is NSF-capable.

```
show cef state
```

```
CEF Status:
 RP instance
 common CEF enabled
IPv4 CEF Status:
 CEF enabled/running
 dCEF enabled/running
 CEF switching enabled/running
 universal per-destination load sharing algorithm, id 7E0E20AE
RRP state:
 I am standby RRP:                   no
 RF Peer Presence:                   yes
 RF Peer Comm reached:               yes
 RF Peer Config done:                yes
 RF Progression blocked:            unblocked (blocked for 00:00:00.588)

 Redundancy mode:                    sso(3)
 CEF NSF sync:                       enabled/running

CEF ISSU Status:
  FIBHWIDB broker
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  FIBIDB broker
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  FIBHWIDB Subblock broker
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  FIBIDB Subblock broker
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  Adjacency update
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  IPv4 table broker
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
  CEF push
    Slot(s): 3 5 40 (0x10000000028) (grp 0x37003204) - Not ISSU aware.
```

Configure OSPF for NSF

The following example shows how to configure OSPF for NSF on a networking device.

```
Router# configure terminal
Router(config)# router ospf 400
Router(config-router)# nsf
```

Verify OSPF for NSF

To verify OSPF for NSF, check whether the NSF is configured on the SSO enabled networking device. Verify that "nsf" appears in the OSPF configuration of the SSO-enabled device by entering the **show running-config** command:

```
Router# show running-config
```

```
router ospf 120
log-adjacency-changes
nsf
network 192.168.20.0 0.0.0.255 area 0
network 192.168.30.0 0.0.0.255 area 1
network 192.168.40.0 0.0.0.255 area 2
```

Next, use the **show ip ospf** command to verify that NSF is enabled on the device.

```
Router# show ip ospf
```

```
Routing Process "ospf 1" with ID 192.168.2.1 and Domain ID 0.0.0.1
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x0
Number of opaque AS LSA 0. Checksum Sum 0x0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled, last NSF restart 00:02:06 ago (took 44 secs)
Area BACKBONE(0)
Number of interfaces in this area is 1 (0 loopback)
Area has no authentication
SPF algorithm executed 3 times
```

DLP-J53 Troubleshoot NSF Using Cisco IOS Commands

Purpose	This procedure troubleshoots NSF using Cisco IOS commands.
Tools/Equipment	None
Prerequisite Procedures	None
Required/As Needed	As needed
Onsite/Remote	Onsite or remote
Security Level	Provisioning or higher

The following commands do not have to be entered in any specific order.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**debug ospf nsf** [**detail**] **Example:** Router# debug ospf nsf [detail]	Displays debugging messages related to OSPF NSF commands.
Step 3	**show cef nsf** **Example:** Router# show cef nsf	Displays the current NSF state of Cisco Express Forwarding on both the active and standby fabric cards.
Step 4	**show cef state** **Example:** Router# show cef state	Displays the Cisco Express Forwarding state on a networking device.
Step 5	**show ip cef** **Example:** Router# show ip cef	Displays entries in the Forwarding Information Base (FIB) that are unresolved, or displays a FIB summary.
Step 6	**show ip ospf** **Example:** Router# show ip ospf	Displays general information about OSPF routing processes.
Step 7	**show ip ospf neighbor** [**detail**] **Example:** Router# show ip ospf neighbor [detail]	Displays OSPF neighbor information on a per-interface basis.
Step 8	**show ip protocols** **Example:** Router# show ip protocols	Displays the parameters and current state of the active routing protocol process.
Step 9	Return to your originating procedure (NTP).	—

In-Service Software Upgrade

Software upgrade is an important consideration for high availability. CPT supports the In-Service Software Upgrade (ISSU) process to perform planned software upgrades within the HA system. ISSU provides the ability to perform a stateful upgrade even when both the fabric cards are of different versions. ISSU is built over the SSO infrastructure.

ISSU allows you to perform a software upgrade or downgrade while the system continues to forward packets. In most networks, planned software upgrades are a significant cause of downtime. ISSU takes advantage of the NSF with SSO and eliminates downtime associated with software upgrades by allowing changes while the system remains in service. ISSU lowers the impact that planned maintenance activities have on network service availability; there is less downtime and better access to critical systems.

The software upgrade is always done at the node level and not directly on the CPT cards. When the package is upgraded, CPT cards must also upgrade to the Cisco IOS image present in the new package. Reverting to the earlier version of the software must be done for all the cards in the shelf and not just for CPT cards.

The Transport Node Controller (TNC) card stores the required images in its flash memory and manages the software upgrade process for all the cards.

SSO and ISSU work together to ensure that the states and configuration of the CPT cards are maintained before and after the upgrade.

Note You can perform the ISSU using CTC and TL1 and not using Cisco IOS commands.

The ISSU message transformations between the active fabric card and the standby fabric card is supported. All the CPT components support the ISSU message transformations between the active fabric card and the standby fabric card.

The ISSU message transformations between the active fabric card and the standby fabric card is supported .The ISSU message transformations between the active fabric card and the line card is also supported. All the CPT components support the ISSU message transformations between the active fabric card and the line card except EVC and QoS components.

Prerequisites for ISSU

- Both the active and standby fabric cards must be available in the system.
- New and existing software images must be loaded on the TNC cards before starting the ISSU process.
- SSO and Cisco NSF must be configured and working properly.

Restrictions for ISSU

- Do not make hardware changes while performing an ISSU process.
- (Recommended) Perform upgrades only during a maintenance window.
- Do not enable new features that require configuration changes, during the ISSU process.

Upgrade Activities in ISSU

The following upgrade activities take place sequentially in a system that has a single fabric card. However, it is recommended that two fabric cards are used in the system.

1. The TNC card boots with the latest image.
2. The TNC card sends the upgrade request to the active fabric card.
3. The active fabric card predownloads its image and the images for the line card.
4. The active fabric card sends the predownload request sequentially to all the line cards based on the slot ID.
5. All the line cards predownload their images.
6. The active fabric card and the line cards reboot with the latest image.

The following upgrade activities take place sequentially in a system that has two fabric cards:

1. The TNC card boots with the latest image.
2. The TNC card sends the upgrade request to both the active and standby fabric cards.
3. The active and standby fabric cards predownload their images and the images for the line card.
4. The active fabric card sends the predownload request sequentially to all the line cards based on the slot ID.
5. All the line cards predownload their images.
6. The active fabric card sends the reboot request to the standby fabric card.
7. The standby fabric card reboots with the latest image.
8. When the standby fabric card reaches the hot standby state, the active fabric card reboots.
9. On upgrade to release 9.3.x, when the rebooted active fabric card joins as standby fabric card, the new active fabric card reloads all the line cards sequentially based on the slot ID.
10. When all the line cards are up, the active fabric card clears the Upgrade alarm on the TNC card.

Table 1: Line Card Reload During Upgrade

Source Release	Destination Release	Line Card Reload
9.5.x	9.5.x	All the line cards reload sequentially.
9.5.x	9.7	All the line cards reload sequentially.

Table 2: Line Card Reload During Downgrade

Source Release	Destination Release	Line Card Reload
9.7	9.5.x	All the line cards reload sequentially.
9.5.x	9.5.x	All the line cards reload sequentially.

Approximate ISSU Duration

CPT supports up to 2 fabric cards, 20 CPT 50 panels, and 4 line cards on a CPT 600 shelf.

Total upgrade time = Boot of TNC card with latest image + Fabric and line card image download on the fabric card + (24 * Predownload of line cards) + Boot of the standby fabric card + Boot of the active fabric card + (4 * Boot of line cards)

24=4 Line cards+16 CPT 50 panels in line card + 4 CPT 50 panels in 2 Fabric cards

Total Upgrade time = 3 minutes + 3 minutes + (24 * 30 seconds) + 6 minutes + 6 minutes + (4 * 5 minutes) = 50 minutes

The above upgrade duration is calculated based on average boot delays. The delays due to ROM MONitoring (ROMMON) and Field Programmable Gate Array (FPGA) upgrades are not accounted in this calculation.

Card Crash During ISSU

When the active fabric card crashes during the predownload of the line card, the standby fabric card becomes active and resumes the predownload process. There are two possible scenarios:

- If the crashed active fabric card recovers and reaches the hot-standby state, the fabric card reboots all the cards sequentially.
- If the crashed active fabric card does not recover, all the cards reboot at once resulting in loss of traffic during the boot duration.

ISSU Commands

The following commands display the ISSU status details of a client.

- **show issu capability**
- **show issu comp-matrix**
- **show issu endpoints**
- **show issu clients**
- **show issu sessions**
- **show issu entities**
- **show issu fsm**
- **show issu negotiated**
- **show issu message**

For more information on these commands, see the *Cisco CPT Command Reference Guide*.

Graceful Restart

LDP graceful restart protects traffic when an LDP session is lost. If an interface that supports a graceful-restart-enabled LDP session fails, MPLS LDP-IGP synchronization is still achieved on the interface while it is protected by graceful restart.

LDP graceful restart must be enabled for LDP to be HA compliant. Graceful restart helps to preserve the MPLS forwarding table entries built by LDP over a SSO.

LDP graceful restart must be enabled before establishing a LDP session. You can configure graceful restart through both CTC and Cisco IOS commands. See DLP-J134 Configure MPLS LDP Graceful Restart Using Cisco IOS Commands and DLP-J135 Configure MPLS LDP Graceful Restart Using CTC for more information.