

Supervisor Redundancy for the Cisco Catalyst 6500 Series Switches with Cisco Catalyst Operating System

Version 3.0



OVERVIEW

Cisco® Catalyst® 6500 Series multilayer switches have become a primary component of an efficient network design in today's enterprise and service provider environments. In such a critical role, the Cisco Catalyst 6500 Series must provide a reliable switching platform that offers high performance and intelligent network services. This paper discusses how the Cisco Catalyst 6500 Series provides high system availability through hardware and software redundancy features. It focuses specifically on the following two areas:

- Supervisor engine redundancy with the Cisco Catalyst Operating System (Catalyst OS) High Availability features, which include the stateful protocol redundancy and image versioning functions
- Cisco IOS® Software Multilayer Switch Feature Card (MSFC) redundancy features—Dual Router Mode (DRM), Configuration-Synchronization (config-sync), Single Router Mode (SRM), and Non-Stop Forwarding with Stateful Switchover (NSF/SSO).

This paper discusses the hybrid software model for the Cisco Catalyst 6500 Series (Catalyst OS on the supervisor engine, Cisco IOS Software on the MSFC) and not on the Cisco IOS Software model (native Cisco IOS Software). For more information about the native Cisco IOS NSF/SSO on the Cisco Catalyst 6500 Series, please visit:

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801c5cd7.shtml. All feature set references will be specifically described as either a Catalyst OS (supervisor engine) or a Cisco IOS Software (MSFC) feature.

The Catalyst OS High Availability feature was first introduced in the Catalyst OS 5.4 release and is available for the Cisco Catalyst 6500 Series Supervisor Engine 1A, Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720. Support for DRM began in Cisco IOS Software Release 12.0(7)XE1 and is available for Supervisor Engine 1A and Supervisor Engine 2. The MSFC config-sync redundancy feature for DRM is supported in Cisco IOS Software Release 12.1(3a)E4 for both the MSFC and MSFC2. The MSFC SRM feature was first supported with Catalyst OS 6.3.1 and Cisco IOS Software Release 12.1(8)E2 for the MSFC2. SRM has also been supported in 12.2SX releases for Supervisor Engine 2, Supervisor Engine 32 and Supervisor Engine 720 until Cisco IOS Software Release 12.2(18)SXF where both SRM and DRM were replaced by the NSF/SSO feature, NSF/SSO being a platform-independent superset of the SRM functionality.

This paper is the third version. This version includes updated sections on the Supervisor Engine 32, Supervisor Engine 720, and Cisco IOS NSF/SSO functionality.

Although component-level redundancy is very important, a high-availability network design relies on the proper combination of individual system redundancy and overall network redundancy. For more detail on high-availability network designs, refer to the publicly available Solution Reference Network Designs white papers at <http://www.cisco.com/go/srnd>.

REDUNDANT SUPERVISOR ENGINES

The High Availability features on the Cisco Catalyst 6500 Series provide low-impact, stateful switchover between redundant supervisor engines. This feature was first available in Catalyst OS Software Release 5.4.

Supervisor Engine Switchover

Dual supervisor engines provide hardware redundancy for the forwarding intelligence of the Cisco Catalyst 6500 Series. The Cisco Catalyst 6500 Series can support up to two supervisor engines (slots 1 and 2 for the Supervisor Engine 1A and Supervisor Engine 2; and slots 1 and 2 in the Cisco Catalyst 6503 and 6504 chassis, slots 5 and 6 in the Cisco Catalyst 6506 and 6509 chassis, and slot 7 and 8 in the Cisco Catalyst 6513 chassis for the Supervisor Engine 32 and Supervisor Engine 720) with one being the running, or active, supervisor engine and the other being the standby supervisor engine. The active supervisor engine is the first one to go online. This can be confirmed by the “Active” LED on the supervisor or by typing the **show module** command from the console. Both supervisor engines must be the same hardware model (that is, if a Policy Feature Card [PFC] and a MSFC are on a Supervisor Engine 1A in slot 1, then a PFC and MSFC must be on a Supervisor Engine 1A in slot 2 as well, or if a Supervisor Engine 2 is in slot 1, a Supervisor Engine 2 must also be in slot 2. Likewise, the PFCx and MSFCx versions must match on the active and standby Supervisor Engine 720 or Supervisor Engine 32. Both Supervisor Engines 1A and 2 can be used in the Cisco Catalyst 6000 and Catalyst 6500 series. However the Supervisor Engine 720 and 32 can only be used with Cisco Catalyst 6500 Series chassis. In the event that an active supervisor engine is taken offline or fails, the standby supervisor engine takes control of the system.

The two supervisor engines in a redundant supervisor engine configuration have different responsibilities. The active supervisor engine is responsible for controlling the system bus and all line cards. All protocols are running on the active supervisor engine and it performs all packet forwarding. The standby supervisor engine does not communicate with the line cards. It receives packets from the network and populates its forwarding tables with this information, but does not participate in any packet forwarding. The relevant protocols on the system are initialized, but not active, on the standby supervisor engine. The Cisco Catalyst 6500 Series supervisor engines are hot-swappable and the standby supervisor engine can be installed in a running, active system. Also please note that redundant supervisor engines do not perform load sharing. The active supervisor engine is providing the entire packet-forwarding intelligence for the system (N+1 redundancy). If the active supervisor engine fails, the standby supervisor engine can still maintain the same system load.

The standby supervisor engine polls the active supervisor engine through the Ethernet out-of-band channel (EOBC) every 5 to 10 milliseconds to monitor the online status of the active supervisor engine. The active supervisor engine might go offline for a variety of reasons such as hardware failures, system overload conditions, memory corruption issues, removal from chassis, being reset by the operator, or real-time diagnostics-driven supervisor switchover (also known as Generic Online Diagnostics)¹. The standby supervisor engine detects this type of failure and becomes the new active supervisor engine. The Catalyst OS software on the supervisor engine is responsible for restoring the protocols, line cards, and forwarding engines to normal operation. This restoration takes place through a fast switchover or a high-availability switchover.

Supervisor Engine Fast Switchover

Because the Catalyst OS High Availability feature was disabled by default until Cisco Catalyst OS Release 8.5, the alternative is referred to as Fast Switchover. The Fast Switchover feature is the predecessor to the High Availability feature and as such is the supervisor engine switchover mechanism in place when High Availability is disabled or not supported in the software version. This feature reduces the switchover time by skipping some events that would typically take place should a supervisor engine fail. Specifically, the fast switchover

¹ Real-time diagnostics driven switchover is only available for supervisor switchovers and not for MSFC switchovers in the Hybrid model.

mechanism allows each line card to skip the respective software downloads and a portion of the diagnostics, which are normally a part of system re-initialization. The switchover still includes restarting all protocols (Layer 2 and above) as well as resetting all ports. The resulting switchover performance with default settings will take approximately 28 seconds plus the time it takes for the protocols to restart. As an example, a switch with the default time values for the Spanning Tree Protocol took approximately 58 seconds after the fast switchover to begin forwarding traffic again. However, the time to begin forwarding traffic after a fast switchover can be reduced by tuning the switch from the default settings. By enabling Portfast, disabling port channels (PagP), and turning trunking off for ports to which workstations are directly attached, the fast switchover time can be reduced to approximately 10 seconds to begin forwarding again. In a live network environment, these supervisor engine switchover times present a major disruption to network operations.

Supervisor Engine High Availability Feature

The High Availability software feature of Catalyst OS further enhances the Cisco Catalyst 6500 Series hardware redundancy by also providing protocol redundancy. This feature includes two main functions: stateful protocol redundancy and image versioning. The High Availability feature must be enabled through the command-line interface (CLI) for these features to operate.

```
Sup-A> (enable) set system highavailability enable
System high availability enabled.
```

As a general practice with redundant supervisors, it is recommended that the High Availability feature be enabled for normal operation.

Supervisor Engine Stateful Protocol Redundancy

The stateful supervisor engine switchover is when the switchover time from the active to the standby supervisor engine is reduced to less than 3 seconds for return to normal operation. This reduced downtime is achieved by synchronizing many of the Layer 2, Layer 3, and Layer 4 protocols² between the active and standby supervisor engines. This is referred to as maintaining protocol state.

For stateful protocol redundancy between dual supervisor engines, a protocol state database is maintained on each supervisor engine for all protocols and features requiring high-availability support. Most of these protocols are only running on the active supervisor engine. In the event of a high-availability switchover, the new active supervisor can start the protocols from the updated database state, rather than the initialization state. This is how a redundant supervisor system can maintain stateful protocol redundancy and minimal network downtime when the active supervisor goes offline. Table 1 shows some protocols and features for high availability that are supported, compatible, and incompatible.

- **High Availability Supported feature**—High availability is fully supported. The state of the feature is preserved between the active and standby supervisor engines in the protocol database.
- **High Availability Compatible feature**—High availability is not supported for these features. The protocol database for these features is *not* synchronized between supervisor engines. The feature *can* be used if the High Availability feature is enabled. For example, if GARP Multicast Registration Protocol (GMRP) and High Availability were both enabled and a high-availability supervisor engine failover took place, the GMRP would be restarted from the initialization state (non-stateful). The stateful protocol redundancy is still in place for the supported features if a compatible feature is enabled.
- **High Availability Incompatible feature**—High Availability is not supported. The protocol database for these features is not synchronized between supervisor engines. The feature should not be enabled if the High Availability feature is enabled. Incorrect behavior may result so these features are not supported with High Availability enabled.

Important: Do not use these features if a high-availability system is required.

The following table lists some supported and non-supported protocols and features for high availability as of the Cisco Catalyst OS Release 8.5.

² Layer 4 protocols include the Layer 4 information in extended IP access lists.

Table 1 High Availability Feature Support

Supported Features	Compatible Features	Incompatible Features
Common Open Policy Service-Differentiated Service (COPS-DS) and COPS for Provisioning (COPS-PR)	Accelerated Server Load Balancing (ASLB)	Dynamic VLANs
Dynamic Trunk Protocol	Cisco Discovery Protocol	Generic VLAN Registration Protocol (GVRP)
Cisco Express Forwarding and adjacency tables	GMRP	Protocol filtering
Private VLANs	Internet Group Management Protocol (IGMP) snooping	
Router access control lists (ACLs)	Remote Monitoring (RMON)	
Multilayer switching (MLS)	Resource Reservation Protocol (RSVP)	
Port Aggregation Protocol/Link Aggregation Protocol (PAgP/LACP)	Simple Network Management Protocol (SNMP)	
Quality-of-service (QoS) ACLs and policers	Telnet sessions	
Switched Port Analyzer (SPAN)	VTP pruning	
Spanning Tree Protocol, Flexlink	Uplinkfast	
Trunking		
UniDirectional Link Detection (UDLD) protocol		
VLAN ACLs		
VLAN Trunking Protocol (VTP)		
Port Security		
802.1x		
Network Admission Control (NAC)		
DHCP snooping		
Dynamic ARP Inspection		
Port ACLs (PAACL)		
Online Diagnostics		

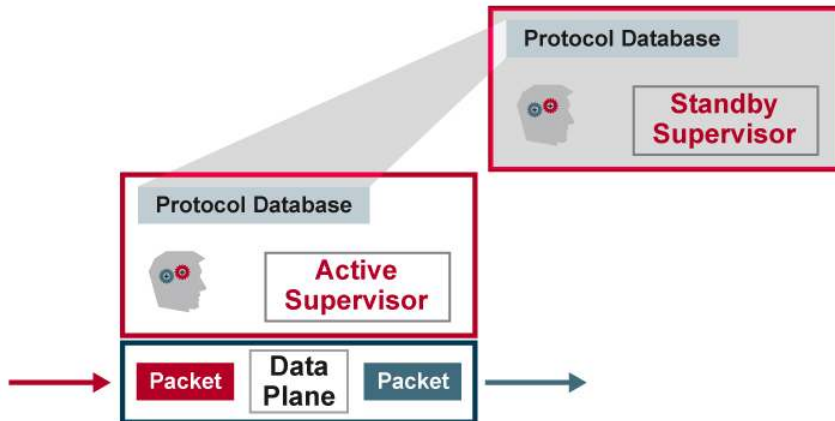
For a current list of the features that are supported with the High Availability feature, see the “Configuring Redundancy” chapter of the Cisco Catalyst 6500 Series Software User Guide and the release notes.

Many Layer 3 and Layer 4 protocols or features are programmed into the application-specific integrated circuits (ASICs) of the PFCx onboard the supervisor engine. Examples of these include access lists (router- and VLAN-based), forwarding tables (multilayer switching cache and Cisco Express Forwarding tables), QoS settings, etc. Upon a supervisor engine failover, these protocols are maintained in the protocol database and will continue to be switched in hardware. Some of these are dependent on a dual MSFC configuration, SRM, or NSF/SSO, which is discussed later in this paper.

The protocol state database, depicted in Figure 1, is a repository of up-to-date protocol state information. It is generated by the active supervisor engine and stored by the standby supervisor engine. The database contains specific system information including module and port states, VLAN information, nonvolatile RAM (NVRAM) configurations, and various protocol-specific data. Both supervisor engines run a synchronizing operation to allow for transfer of this data. When a database entry is updated on the active supervisor engine, the synchronizing operation on the active supervisor engine transfers the updated entry to the standby supervisor engine. The standby

supervisor engine's synchronizing process receives these asynchronous updates and enters them into the protocol state database on the standby supervisor engine. When the system starts or when a second supervisor engine is hot-inserted, a global synchronization takes place between the protocol databases to ensure all protocol states are up to date.

Figure 1. Stateful Protocol Database Depiction



In summary of the stateful protocol feature, the high-availability switchover performance is more dependent on the status of the synchronization procedure than on the complexity of the configuration. After the system and protocols reach a stable operating point, the protocol state databases on each of the supervisors will have a fairly similar status. The resulting high-availability supervisor engine switchover performance is less than 3 seconds.

SUPERVISOR ENGINE SOFTWARE IMAGE UPGRADES

In a redundant supervisor engine configuration, Catalyst OS images need to be properly managed to ensure high availability of the system. The following section describes some options for managing Catalyst OS images.

Supervisor Engine Image Synchronization

By default on the Cisco Catalyst 6500 Series, the Catalyst OS software images on the active and the standby supervisor engines must be the same. This allows the system to maintain a stable operating environment by ensuring that the supervisor engine switchover occurs with the same software features and revisions on the new active supervisor engine as on the previously active supervisor engine. If the two supervisor images are not the same version during system bootup, the active supervisor engine downloads its current boot image to the standby supervisor engine. The NVRAM configuration of the active supervisor engine is also synchronized between the supervisor engines.

The image synchronization feature of Catalyst OS provides software consistency between supervisor engines. It does not, however, allow for software to be upgraded without taking the system offline for an extended period of time. To perform the upgrade, the active supervisor engine requires a supervisor engine reset to load the new version of software. It will then synchronize the software images to the standby supervisor engine. This typically has to be performed during a scheduled downtime or maintenance window because the entire system will need to be warm booted. Also note that the MSFC Cisco IOS Software is not a part of this synchronization process.

Supervisor Engine High-Availability Versioning Feature

Versioning is the second portion of the Catalyst OS High Availability feature and is dependent on having the High Availability feature enabled in a dual supervisor engine configuration. As such, it allows different but compatible images to be running on the active and standby supervisor engines, thus disabling the default supervisor image synchronization process. This feature is primarily used to simplify

the software upgrade process when two supervisor engines are involved. Please note that a stateful switchover will not be possible for upgrades between different images except for a few corner cases.

If two different image versions are running, the system will determine the compatibility of these two versions. The active and standby supervisor engines exchange image version information to determine if the two software images are compatible. Image versions are defined as one of three options: compatible, incompatible, or upgradeable. Compatible versions imply that stateful protocol redundancy can be supported between the different images. All configuration settings made to the NVRAM on the active supervisor can be sent to the standby supervisor engine. Two Catalyst OS versions are incompatible if synchronizing the protocol state databases between the two versions is not possible. If two software images are incompatible, the software upgrade process will impact the system operation (that is, be greater than the one to three second switchover time of a high-availability switchover) and no NVRAM configuration changes will be synchronized between supervisor engines. A special case of incompatible versions is referred to as upgradeable. This case means that the high-availability supervisor engine switchover is not available, but configuration changes to the NVRAM on the active supervisor engine can be synchronized to the standby supervisor engine. This is more of a special case because it allows two different software versions to run with synchronized configurations but without the ability for a failover.

If the Catalyst OS software images are anything but compatible, the high-availability switchover will not be possible. The operational status output from the command, **show system highavailability** should be monitored to determine the high-availability compatibility of two Catalyst OS images. The operational status can either be ON or OFF (with some system specific status messages). The following output shows that high availability is enabled and that the Catalyst OS versions are high-availability-compatible (Op-status: ON).

```
Sup-A> (enable) show system highavailability
Highavailability: enabled
Highavailability versioning: disabled
Highavailability Operational-status: ON
```

As a general practice, it is recommended that high-availability versioning be enabled only when upgrading the Catalyst OS software. The traditional image synchronization process (high-availability versioning disabled) should be implemented for normal operating conditions. Generally speaking, high-availability compatible images are only available between maintenance releases of the Catalyst OS software. A maintenance release is a new version of software with incremental feature upgrades and bug fixes such as upgrading from version 5.5.1 to version 5.5.2. Major releases will not be high-availability compatible. The release notes include a high-availability compatibility listing at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/ol_4498.htm#wp292496.

CATALYST OS IMAGE UPGRADE PROCEDURE

Based on the discussion above, the following is provided as a recommendation for how to perform a software upgrade with redundant supervisor engines. While the MSFC is affected by this procedure, it is not discussed in this procedure. A discussion is included in the section "MSFC High Availability Features."

In this example, the supervisor engine in slot 1 (Sup-A) will begin as the active supervisor engine and the supervisor engine in slot 2 (Sup-B) will begin in standby mode. It is recommended that a console connection be available for both supervisor engines for this procedure.

1. Disable the High Availability feature on the active supervisor engine.

```
Sup-A> (enable) set system highavailability disable

This feature is disabled by default.
```

2. Load the new Catalyst OS software image into the boot Flash (via slot0, Trivial File Transfer Protocol [TFTP], etc.) of only the active supervisor engine.

```
Sup-A> (enable) copy slot0:cat6000-sup2k8.7-2-2.bin bootflash:cat6000-sup2k8.7-2-2.bin
```

3. Verify that the new image is located in the bootflash of the active supervisor engine.

```
Sup-A> (enable) dir bootflash:
```

4. Clear the current boot variable.

```
Sup-A> (enable) clear boot system all
```

5. Set the boot variable on the active supervisor engine to the new Catalyst OS software image.

```
Sup-A> (enable) set boot system flash bootflash:cat6000-sup2k8.7-2-2.bin
```

In approximately 120 seconds, the image set as the boot entry on the active supervisor engine will be copied to the bootflash on the standby supervisor engine (this is the image synchronization). This is an internal TFTP of the Catalyst OS image file and takes a few minutes to complete. The image file will have a BTSYNC appended to the beginning of the filename. This is to designate that it has been synchronized from the active supervisor engine's boot-time image.

6. When the images have been synchronized, verify that the new image is located on the standby supervisor engine and the boot variable is properly set.

```
Sup-A> (enable) dir 2/bootflash:
```

```
Sup-A> (enable) show boot 2
```

The new Catalyst OS image is now on both supervisor engines.

7. Enable high-availability versioning on the active supervisor engine.

```
Sup-A> (enable) set system highavailability enable
```

```
Sup-A> (enable) set system highavailability versioning enable
```

Before the standby supervisor engine running the new software becomes active, versioning must be enabled. This allows the standby supervisor engine to reboot under the new version of Catalyst OS while remaining the standby supervisor engine.

8. It is the intent of these upgrade procedures to allow for a fallback plan of using the old Catalyst OS image. The now-active supervisor engine must maintain that older image (even after an accidental reboot). Therefore, the boot variable on the active supervisor engine must be changed to its original setting, which should still be stored in the bootflash.

```
Sup-A> (enable) set boot system flash bootflash:cat6000-sup2k8.old.bin
```

Note: Since versioning is enabled, the setting of the boot variable does not cause an image synchronization.

9. Reset the standby supervisor engine.

```
Sup-A> (enable) reset 2
```

The standby supervisor engine will reboot with the new Catalyst OS image. It will remain the standby supervisor engine and will not affect the operation of the active supervisor engine.

10. After the standby supervisor engine has rebooted, verify that it is running the new Catalyst OS image.

```
Sup-A> (enable) show module
```

The standby supervisor engine should show the new software version and it should be different from the active supervisor engine's software version.

11. Verify that the two different Catalyst OS images are high-availability compatible.

```
Sup-A> (enable) show system highavailability
```

For the high-availability switchover to occur, the operational status of the High Availability feature must be ON. If it is not, the system will be upgraded with a fast switchover (nonstateful) and the protocols will need to be restarted.

12. Reset the active supervisor engine. You will need to change the console connection to the supervisor engine in slot 2 (Sup-B) to maintain command-line operation.

```
Sup-A> (enable) reset 1
```

The standby supervisor engine now takes over as the active supervisor engine (running new software) and the previously active supervisor engine is rebooted to become the new standby supervisor engine. This supervisor engine failover will disrupt a slight amount of traffic through the device; the amount of traffic that is affected depends on whether a high-availability switchover or a fast switchover takes place.

13. Verify the system is performing as expected. The supervisor engine in slot 2 is now the active supervisor engine running the new version of the Catalyst OS software. The supervisor engine in slot 1 is now the standby supervisor engine running the previous version of software. As such, the new standby supervisor engine can be used as a fallback plan to revert to the previous version of Catalyst OS.

14. If the system is operating as expected, the boot configuration on the standby supervisor engine (now Sup-A) will need to be updated.

This can be accomplished by disabling versioning on the new active supervisor engine, which automatically enables the image synchronization feature.

```
Sup-B> (enable) set system highavailability versioning disable
```

```
Sup-B> (enable) reset 1
```

The supervisor Catalyst OS software upgrade procedure is now complete.

MSFC HIGH AVAILABILITY FEATURES

The MSFC routing engine is a daughter card in the supervisor engine and is available in four versions: MSFC, MSFC2, MSFC2a, and MSFC3 (see MSFC datasheet for configuration requirements). The MSFC is optional on the Supervisor Engine 1A and Supervisor Engine 2 but it is built in on the Supervisor Engine 32 and Supervisor Engine 720. A redundant supervisor engine hardware configuration can also include redundant MSFC routing engines. As such, the proper operation of the MSFC is predicated by proper operation of the supervisor engine. From a switchover operation perspective, the Cisco IOS MSFC switchover is independent of the supervisor engine Catalyst OS

high-availability switchover, allowing a model where the active supervisor engine may be in slot 1 and the active MSFC may be in slot 2. However, a supervisor engine reset or failover causes systematically an MSFC reset or failover.

While the Catalyst OS High Availability feature maintains the protocol state between redundant supervisor engines, the dual MSFC operation operates under three different redundancy modes: DRM, SRM, and NSF/SSO. When running either MSFC redundancy mode, Cisco Systems® recommends that the Catalyst OS High Availability feature be enabled.

Dual Router Mode

DRM is the original MSFC configuration for redundant supervisor or MSFC configurations. DRM is supported for Supervisor Engine 1A and Supervisor Engine 2 only. It is not supported with Supervisor Engine 32 and Supervisor Engine 720 and was replaced by the NSF/SSO redundancy mechanism in Cisco IOS Software Release 12.2(18)SXF. In this mode, both MSFCs are active routers on the network. Having two active MSFCs in a single chassis does not mean having two separate routers. In fact, both MSFCs must have a nearly identical configuration, as described below in more detail. The main idea for DRM is that each MSFC independently builds an accurate picture of the Layer 3 network.

DRM Operation

The failover mechanism between MSFCs in DRM is the Hot Standby Routing Protocol (HSRP). HSRP allows the two MSFCs to maintain internal communication and react to an MSFC failover. HSRP needs to be configured on both MSFCs for each VLAN where first hop default gateway redundancy is required. Internal HSRP between MSFCs works in the same manner as HSRP between physically separate devices by sending hello messages between the routing engines. For more information about configuring HSRP, see the Cisco IOS Software Configuration Guides at:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdip.htm#wp1001531 and

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/redund.htm#wp1058202.

Because both MSFCs have independent routing tables, there is little routing protocol convergence necessary in the event of an MSFC failure. With DRM and based on HSRP timers, the MSFC failover can be configured to less than 3 seconds for LAN interfaces, thus aligning the Layer 3 failover of the MSFC with the supervisor engine failover time.

Because each MSFC has the potential for taking over for the other one, they need to maintain identical configurations. This is an extremely important point to understand in DRM. Configuration parameters such as interfaces, access lists, policy routing, etc. must be configured exactly the same on both MSFCs. Parameters that cannot be duplicated on a network such as IP addresses and HSRP settings are the only parameters that are configured differently on each MSFC.

The MSFC is responsible for programming certain functions of the ASIC hardware on the PFCx. The first MSFC to go online is considered the *designated router* and the second MSFC is considered the *nondesignated router*. In a Supervisor Engine 1A system, both the designated router and the nondesignated router are able to program Layer 3 entries into the PFC NetFlow table for routing functions. In a Supervisor Engine 2 system, only the designated router programs the Layer 3 entries in the PFC2s Cisco Express Forwarding table. For both Supervisor Engines 1A and 2, all router ACLs and multicast shortcuts are programmed from the designated router. As you can see, the requirement for each MSFC to have an identical configuration is a necessity. If the MSFCs in DRM have different configurations, the forwarding ASICs will be programmed incorrectly, resulting in unexpected behavior.

MSFC Configuration Synchronization

Beginning with Cisco IOS Software Release 12.1(3a)E4 on the MSFC, an MSFC redundancy feature called config-sync has been available to streamline the redundant MSFC configuration process for both MSFC and MSFC2. This feature can be used to simplify configuration of the two MSFCs and to ensure that the MSFC configurations match. Both the startup and running configurations between the designated (primary) and nondesignated (secondary) MSFCs are synchronized. Specifically, when a **write memory** or **copy <source> startup-config** command is issued on the designated MSFC, the startup configurations in NVRAM of both MSFCs are updated. This allows the

configurations on the designated and nondesignated MSFCs to maintain the same configuration without having to manually type each command twice.

The following commands enable MSFC config-sync:

```
MSFC-Sup-15 (config)# redundancy
MSFC-Sup-15 (config-r)# high-availability
MSFC-Sup-15 (config-r-ha)# config-sync
```

With config-sync, all configurations for the designated and nondesignated MSFCs are done through the CLI of the designated MSFC. Configuration of the nondesignated MSFC is accomplished through the use of the alt keyword. This is the only way to configure the nondesignated MSFC when config-sync is enabled. For example:

```
MSFC-Sup-15 (config-if)# ip address a.b.c.1 x.x.x.0 alt ip address a.b.c.2 x.x.x.0
MSFC-Sup-15 (config-if)# standby 10 priority 100 alt standby 10 priority 50
```

The command syntax does not change. The portion of the command listed before the **alt** keyword applies to the MSFC in slot 1 and the portion of the command listed after the **alt** keyword applies to the MSFC in slot 2. The config-sync feature is only supported for general IP or IPX configurations; configuration parameters for Appletalk, DECnet, etc., do not have corresponding **alt** keyword options.

WAN Interfaces in DRM

In DRM, the interfaces of a WAN module (either an Optical Service Module [OSM] or FlexWAN) are managed by only the designated MSFC. Prior to enabling the config-sync feature, the WAN interfaces do not show up in the nondesignated MSFC configuration, so they are not configurable on the nondesignated MSFC. During a supervisor engine or MSFC failover, the MSFC that becomes the new designated MSFC will not have properly configured WAN interfaces. For this reason a redundant supervisor engine or MSFC configuration without config-sync was not supported with WAN modules installed. By enabling the MSFC config-sync feature, this limitation is removed and WAN modules are supported in a redundant supervisor engine configuration. WAN modules should not reset during a high-availability switchover with config-sync enabled.

DRM Challenges

DRM was the original option for MSFC redundancy. This solution has been very successful by allowing for stateful Layer 3 failover between MSFCs, but it also introduces some complexity into network design and administration. The following three points present scenarios where DRM does not provide the best solution for Layer 3 redundancy:

- Each MSFC must have a unique IP address for each VLAN interface. In a distribution or core implementation using DRM as well as dual chassis, this could require up to five router IP addresses to be allocated per VLAN (four router addresses plus one HSRP address). This also increases the number of routing protocol neighbors, which can add to the CPU burden on a router. The tasks of addressing and managing four routers in this case can be a challenge that outweighs the benefits of added redundancy.
- In a redundant configuration where multiple MSFCs are connected to the same Ethernet segment, only one MSFC forwards the multicast traffic from the source to the receivers on the outgoing interfaces. The Protocol Independent Multicast designated forwarder (PIM-DF) forwards the data in the common VLAN, but the non-PIM-DF receives the forwarded multicast traffic as well. The redundant MSFC (non-PIM-DF) must drop this traffic because it has arrived on the wrong interface and will fail the reverse path forwarding (RPF) check. Traffic that fails the RPF check is called non-RPF traffic. In general, routers may not manage non-RPF traffic efficiently. With DRM, there is at least one router (the other MSFC) on each VLAN that will receive this non-RPF traffic.
- The requirement for exact configuration parameters on both MSFCs has been a complicated point for many customers. The effort to ensure that all configuration parameters are the same is a challenge when working with large Cisco IOS configuration files. Feature enhancements such as config-sync have been developed to simplify this process but do not scale.

For these scenarios, SRM and NSF/SSO are now available.

Single Router Mode

Single Router Mode (SRM) is provided as an option for customers who wish to implement redundant supervisor engines or MSFCs in a system with only one active router in a chassis. SRM has the ability to use the Layer 2 and Layer 4 redundancy of Catalyst OS on the supervisor engine as well as a streamlined approach to Layer 3 redundancy. The minimum software requirements are Catalyst OS 6.3.1 and Cisco IOS Software Release 12.1(8)E2 for the MSFC. SRM has been supported on Supervisor Engine 1A, Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720. SRM has been replaced by NSF/SSO starting with Cisco IOS Software Release 12.2(18)SXF.

SRM improves upon the challenges faced with DRM (described earlier). Specifically, SRM provides the following:

- A reduction in Layer 3 complexity for IP addressing and routing protocol neighbor relationships.
- A fix for the non-RPF traffic issue with having two active multicast routers on the same segment (because there is only one active router in the chassis with SRM).
- A simpler configuration for the user because only a single command set is entered from one CLI and it applies to the active router. This eliminates the challenge of ensuring that both MSFCs have the same configurations.

The following commands enable SRM:

```
MSFC-Sup-15 (config)# redundancy
MSFC-Sup-15 (config-r)# high-availability
MSFC-Sup-15 (config-r-ha)# single-router-mode
```

SRM Operation

In this mode, only the designated router will be visible to the network at any given time. The nondesignated router will be started and will maintain exactly the same configuration as the designated router (the configurations are automatically synchronized when SRM is active). However, the nondesignated router's interfaces will be kept in a line down state and not visible to the network. Routing protocol processes are also created on the nondesignated router, but they do not send or receive updates from the network because all the interfaces are down. This is verified from the Catalyst OS command line below; note both the supervisor and the MSFC in slot 2 are listed as standby.

```
SRM> (enable) show module
Mod Slot Ports Module-Type           Model                Sub Status
-----
 1   1   2   1000BaseX Supervisor      WS-X6K-SUP2-2GE     yes ok
15   1   1   Multilayer Switch Feature WS-F6K-MSFC2        no  ok
 2   2   2   1000BaseX Supervisor      WS-X6K-SUP2-2GE     yes standby
16   2   1   Multilayer Switch Feature WS-F6K-MSFC2        no  standby
```

If the designated router fails in an SRM configuration, the other MSFC changes state from nondesignated router to designated router. This new designated router changes its interface state to link up and begins to build its routing table. It follows that the control plane failover time will be proportional to the routing protocol configuration and complexity. However, there are existing Layer 3 forwarding entries in the PFCx which are used to forward routed traffic in the hardware path. The high-availability functions of the Catalyst OS are used to maintain this forwarding information after a failover. This allows for minimal impact to the Layer 3 data plane traffic while the Layer 3 control plane converges. After the MSFC builds its routing table, the entries in the PFCx can be updated.

Beginning in Cisco IOS Software Release 12.1(11b)E, there is a transition timer feature for running SRM on the Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720³. This timer configures the time that the new designated router will wait before

³This feature does not apply to the supervisor1a/PFC/MFSCx because the PFC uses a flow-based forwarding architecture and all new flows are initially sent to the MSFCx software path.

downloading any new hardware Cisco Express Forwarding entries to the PFCx. Because of differences in routing convergence times, the default of 120 seconds might not be long enough to allow for complete convergence before programming the PFCx hardware.

The same IP and MAC addresses are used for the designated router, whether or not the MSFC is the designated router. The MSFC that is chosen as the designated router will communicate its default MAC address to the MSFC that is the nondesignated router. All subsequent interfaces created on the nondesignated router use this MAC address, unless the user explicitly configures a different MAC address.

On bootup the two MSFCs perform a “handshake” process, which takes about a minute, before entering SRM mode. Do not make configuration changes on the nondesignated router during this period.

The following commands can be used to verify that SRM is enabled:

```
SRM# show redundancy
Designated Router: 1 Non-designated Router: 2
Redundancy Status: designated
Config Sync AdminStatus: enabled
Config Sync RuntimeStatus: enabled
Single Router Mode AdminStatus: enabled
Single Router Mode RuntimeStatus: enabled
Single Router Mode transition timer: 120 seconds
```

For more details on configuring SRM, see section “MSFC Redundancy - Single Router Mode Redundancy” in the Catalyst OS configuration guide at: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/redund.htm.

WAN Interfaces in SRM

Because the MSFC configurations are synchronized as an inherent part of SRM, WAN modules (OSMs and FlexWAN) are supported with redundant supervisor engines or MSFCs configured for SRM. As in DRM, the designated router manages the WAN interfaces. The interfaces are configured completely on the designated router and that configuration is synchronized to the nondesignated router. For failover scenarios, the new designated router will take ownership of the WAN interfaces as soon as that MSFC becomes the designated router. Additionally, the WAN modules should not reload upon a high-availability switchover. With SRM enabled, there is no manual configuration necessary on the WAN interfaces to support an MSFC failover.

SRM Configuration and Conversion Procedure

The configuration guide has a very good procedure for configuring SRM, converting from DRM to SRM, and performing software upgrades with SRM enabled. The latest recommended procedures are available at the following URLs:

- "Configuring Single Router Mode Redundancy"
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/redund.htm.
- "Upgrading Images with Single Router Mode Enabled"
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_4/config_gd/redund.htm.

SRM and IP Multicast

With SRM, the VLAN interfaces on the nondesignated router are in a *down* state. Even after a failover, these interfaces will not move into an *up* state until the supervisor engine verifies that the VLAN has at least one connected physical interface in the forwarding state. This interruption causes the supervisor engine to delete all multicast entries in the PFCx and this disrupts multicast forwarding. As an enhancement to the original SRM implementation, Catalyst OS release 7.1 provides support for IP Multicast stateful redundancy (Multicast MLS). When SRM is enabled in Catalyst OS release 7.1, the multicast flows are preserved during a failover.

Non-Stop Forwarding with Stateful Switchover

Starting with Cisco IOS Software Release 12.2(18)SXF, NSF/SSO supersedes the SRM redundancy mode. Just like SRM with High Availability, NSF/SSO provides Stateful Switchover at Layer 2 and Layer 4. Layer 3 interfaces do not bounce upon switchover, hardware tables are synchronized between the active and standby supervisor engines, and configuration changes made on the active MSFC are reflected on the standby MSFC. NSF/SSO is supported on the Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720. To obtain the full NSF/SSO functionality, the switch processor must be upgraded to Catalyst OS software 8.5(1) and later releases.

SRM is no longer an option in Cisco IOS Software Release 12.2(18)SXF because it was replaced by NSF/SSO, which provides a superset of the SRM functionality. Along with NSF/SSO, Release 12.2(18)SXF supports the Route Processor Redundancy (RPR) mode. RPR mode does not provide stateful MSFC redundancy; only minimal information such as running and startup configuration are synchronized between the active and the standby MSFC, routed interfaces flap, and routing processes need to be restarted upon MSFC failover. However, RPR allows different Cisco IOS Software versions to coexist on the active and standby MSFC, which is helpful during software upgrades. But it is recommended to always configure the MSFC redundancy mode to NSF/SSO to achieve the highest availability.

NSF/SSO improves upon the SRM Layer 3 functionality. Specifically, NSF/SSO provides the following:

- A platform-independent redundancy mechanism available across Cisco platforms, allowing ease of configuration and troubleshooting
- Routing protocol stability and continuous Layer 3 forwarding upon switchover: An NSF/SSO switchover is not associated with a route reconvergence event or routing flaps on NSF-aware neighbor routers. Instead, while SRM also ensures continuous forwarding in hardware upon switchover, SRM does not implement the NSF routing protocol extensions that prevent routing flaps as a result of neighbor peering relationships being torn down upon failover.

The following commands enable NSF/SSO on the MSFC:

```
MSFC (config)# redundancy  
MSFC (config-r)# mode sso
```

Configuring the redundancy mode to **mode sso** provides the equivalent to the SRM functionality and should be done automatically when upgrading from an SRM software version to Cisco IOS Software Release 12.2(18)SXF or later. Furthermore, a routing-protocol-specific configuration can optionally be used to enable NSF for a given routing protocol. NSF is supported as of Cisco IOS Software Release 12.2(18)SXF for Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS) Protocol, and Border Gateway Protocol (BGP) and can easily be enabled with the **nsf** keyword for EIGRP, OSPF, and IS-IS and by enabling the **graceful-restart** capability for BGP under the routing process. The following command enables NSF for EIGRP:

```
MSFC (config)# router eigrp 10  
MSFC (config-router)# nsf
```

The following document provides detailed NSF/SSO configuration and the most up-to-date NSF/SSO information:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/nsfss0.htm#wp1119658>.

NSF/SSO Operation

NSF/SSO works in conjunction with the Catalyst OS High Availability feature to help ensure Layer 3 integrity following a switchover. It allows a router experiencing the failure of an active supervisor engine to continue forwarding data packets along known routes while the routing protocol information is recovered and validated. This forwarding can continue to occur even though peering arrangements with neighbor routers have been lost on the restarting router. NSF/SSO relies on the separation of the control plane and the data plane during supervisor engine switchover. The data plane continues to forward packets based on pre-switchover Cisco Express Forwarding information. The control plane implements NSF routing protocol extensions (sometimes referred as graceful restart extensions) to signal a supervisor engine restart to NSF-aware neighbor routers, reform its neighbor adjacencies, and rebuild its routing protocol database transparently following a switchover.

An *NSF-capable* router implements the NSF functionality and continues to forward data packets after a supervisor engine failure. An NSF-aware router is a router running a software version that understands the NSF routing protocol extensions: An *NSF-aware* router does not tear down its neighbor relationships with the NSF-capable restarting router, and it can help a neighboring NSF-capable router gracefully restart, thus avoiding unnecessary route flaps and network instability. An NSF-capable router is also NSF-aware. On the Cisco Catalyst 6500 Series, routers can be made NSF-aware by upgrading the supervisor engine Cisco IOS Software to Release 12.2(18)SXD and later or by upgrading the hybrid MSFC Cisco IOS Software to Release 12.2(18)SXF and later.

The NSF/SSO operation on the Cisco Catalyst 6500 Series is fully described in the white paper available at:

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801c5cd7.shtml.

WAN Interfaces in NSF/SSO

Because the MSFC configurations are synchronized as an inherent part of NSF/SSO, all WAN modules are supported with redundant supervisor engines or MSFCs configured for NSF/SSO. As in SRM, the designated router manages the WAN interfaces. The interfaces are configured completely on the designated router and that configuration is synchronized to the nondesignated router. For failover scenarios, the new designated router will take ownership of the WAN interfaces as soon as that MSFC becomes the designated router. Additionally, the WAN modules should not reload upon a high-availability switchover when NSF/SSO is enabled. With NSF/SSO enabled, there is no manual configuration necessary on the WAN interfaces to support an MSFC failover. However, WAN module interfaces go down and then come back up during a stateful switchover, and routing protocols do not perform NSF if NSF is configured over WAN interfaces.

NSF/SSO Configuration, Conversion Procedure, and Considerations

Fast Software Upgrade (FSU) can be used to minimize downtime associated with a planned MSFC software upgrade. With this process, the redundancy mode reverts transparently to RPR during the upgrade and will transparently be set back to SSO when the active and standby MSFC image versions match. The procedure impacts the data traffic and it is thus recommended to perform this procedure during a scheduled maintenance window.

The procedure is described as follows: In order to run in SSO redundancy mode, Cisco IOS image versions must be the same on the redundant and active MSFC. In this redundancy mode, the active MSFC checks the Cisco IOS image version of the redundant MSFC when the redundant MSFC comes online. If the Cisco IOS image version on the redundant supervisor engine does not match the Cisco IOS image version on the active supervisor engine, the software sets the redundancy mode to RPR while doing a software upgrade and sets it back to SSO when the software upgrade is complete. The MSFC FSU operation is described in the following document:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_5/config_gd/nsf_sso.htm#wp1064945.

The SRM-to-NSF/SSO conversion procedure is as follows: An image upgrade from a Cisco IOS pre-12.2(18)SXF release to a Cisco IOS post-12.2(18)SXF MSFC image will automatically force the redundancy mode from SRM to NSF/SSO because SRM is no longer available in post-12.2(18)SXF releases. As such, the conversion procedure from SRM to NSF/SSO does not require any configuration. However, the FSU procedure does not apply when converting from SRM to NSF/SSO; in order to upgrade MSFC images from a Cisco IOS pre-12.2(18)SXF release to a Cisco IOS post-12.2(18)SXF MSFC image, an administrator is required to load the new images on each of the MSFCs and then simultaneously boot the MSFCs. During this time, the MSFCs will be offline.

Additional configuration can be entered to provide the Layer 3 NSF capability, which is enabled on a per-protocol basis under the routing process. Note that NSF/SSO replaces SRM and that as such, the SRM transition time (specifying the amount of time the newly designated router waits before downloading the new Layer 3 switching information to the switch processor) is no longer needed nor an option. NSF configuration steps can be found at: <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/nsfssso.htm>.

In addition, the supervisor engine should be upgraded to Catalyst OS Software Version 8.5.

In order to fully benefit from NSF and avoid routing protocol flaps upon NSF/SSO switchover, a few considerations need to be taken into account: neighbors need to be NSF-aware, and Interior Gateway Protocol (IGP) timers must not be tuned too aggressively. Tuning IGP

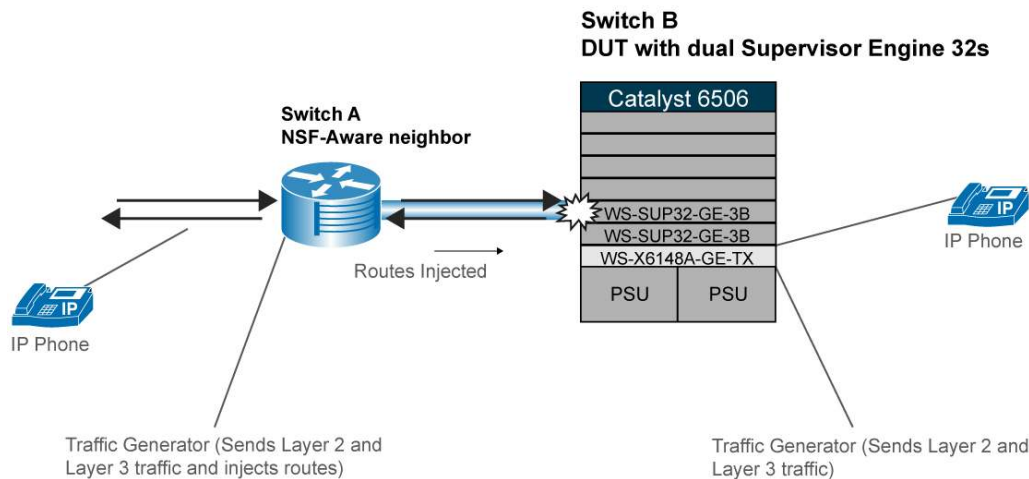
timers too aggressively would negate the graceful restart mechanism because these mechanisms rely on successfully signaling neighbors that a graceful restart/switchover has occurred before the routing protocol dead/hold timers expire.

More information on Cisco NSF deployment and Cisco NSF Timer Manipulation is available at:
http://www.cisco.com/en/US/products/ps6550/prod_white_papers_list.html.

SUPERVISOR ENGINE AND MSFC FAILOVER TESTS

Layer 2 and Layer 3 failover tests were performed on the Cisco Catalyst 6500 Series in order to compare the SRM with High Availability to the NSF/SSO mechanisms performance. Figure 2 represents the test bed that was created for this purpose. The device under test (DUT) has dual Supervisor Engine 32s and the software versions tested include the Catalyst OS 8.5 and Cisco IOS Software 12.2(17d)SXB11 releases for the SRM with High Availability test cases and the Catalyst OS 8.5 and Cisco IOS Software 12.2(18)SXF releases for the NSF/SSO test cases. Tests were initiated by resetting the active MSFC or pulling off the active supervisor engine. Each scenario was tested three times and the results were averaged.

Figure 2. Setup for Layer 3 and Layer 2 Failover Tests



Layer 3 Failover

For the Layer 3 Failover tests, OSPF is configured on the DUT and the switch A. Switch A is running Cisco IOS Software Release 12.2(18)SXF and is NSF-aware. 10000 static routes are injected in the OSPF database on switch A. The traffic flow consists of Layer 3 bidirectional traffic flowing between a traffic generator port attached to switch A and another traffic generator port attached to the DUT.

When failing over the MSFC2a only, the packet loss reported with the NSF/SSO implementation is averaged at 50 milliseconds (ms). The same test leads to 11seconds (s) of average packet loss for Layer 3 traffic flowing from switch A to switch B with the NSF/SSO implementation if NSF is not configured and it also leads to 11s of average packet loss with the SRM/SSO implementation. The bulk of the packet loss comes from the fact that switch A would have torn down its OSPF adjacency with the DUT after switchover, therefore leading to an OSPF convergence event on switch A and to traffic loss even though the DUT still has a valid forwarding path. For traffic flowing from switch B to switch A, traffic loss would be minimal because switch B's forwarding data plane can be used while OSPF converges.

When failing over the supervisor engine (by pulling the active supervisor engine or resetting the Catalyst OS designated supervisor engine), results lead to an average 200-ms packet loss with the NSF/SSO implementation and lead to an average of 11seconds with the SRM/SSO implementation for traffic affected by dynamic routing (traffic flowing from switch A to switch B) convergence events.

The complete MSFC2a configurations are as follows:

SRM

```
hostname SRM
!
redundancy
  high-availability
  single-router-mode
  mode none
!
boot system flash bootflash:c6msfc2a-jsv-mz.122-17d.SXB11
!
interface Vlan5
  description to traffic generator
  ip address 10.5.10.1 255.255.255.0
!
interface Vlan45
  description to switchA
  ip address 10.45.45.5 255.255.255.0
!
router ospf 1
  log-adjacency-changes
  passive-interface default
  no passive-interface Vlan45
  network 10.5.10.0 0.0.0.255 area 0
  network 10.45.45.0 0.0.0.255 area 0
!
```

NSF/SSO

```
hostname Sup32-NSF-SSO
!
redundancy
  mode sso
boot system flash bootflash:c6msfc2a-entservicesk9_wan-mz.122-18.SXF
!
interface Vlan5
  description to traffic generator
  ip address 10.5.10.1 255.255.255.0
!
interface Vlan45
  description to switchA
  ip address 10.45.45.5 255.255.255.0
!
router ospf 1
```

```
log-adjacency-changes
nsf
passive-interface default
no passive-interface Vlan45
network 10.5.10.0 0.0.0.255 area 0
network 10.45.45.0 0.0.0.255 area 0
!
```

Layer 2 Failover

For the Layer 2 failover tests, Layer 2 bidirectional traffic flowed on VLAN10 from the traffic generator port connected to switch A to the traffic generator port connected to switch B. As expected, the SRM with High Availability and the NSF/SSO implementation provide the same failover results when failing over the active supervisor engine to the standby supervisor engine. Indeed, only the Catalyst OS high availability implementation is a factor here and the MSFC2a SRM or NSF implementations do not play a role here. Traffic loss averaged at 200 ms.

The same testbed was used to verify that IP phone calls did not drop following a switchover when established over a Layer 2 and Layer 3 network. An IP phone call was established between one IP phone connected to switch A and one IP phone connected to the DUT. With the Catalyst OS High Availability feature enabled and the NSF/SSO mechanism enabled on the MSFC2a (when the phone call is established over a Layer 3 network), a supervisor engine switchover was initiated. The IP phone call was maintained through the supervisor engine switchover and the call participants hardly noticed the disruption. This provides a real-world example of the Catalyst 6500 Series Switch's ability to provide high availability in all layers of the network.

For general purposes, it is still maintained that the Layer 2 and Layer 3 stateful supervisor engine switchover will take place in less than 3 seconds in most real-world scenarios.

CONCLUSION

The high-availability and redundancy features of the Cisco Catalyst 6500 Series provide a very reliable switching and routing platform. Hardware redundancy in the form of dual supervisor engines, dual routing engines, dual switching fabrics, multiple fans, and dual power supplies help to reduce the potential downtime of a network. Software redundancy features such as the High Availability feature of Catalyst OS and the DRM, SRM, or NSF/SSO options for MSFC failover build on this hardware redundancy for a very stable operating environment. The combination of these features, in addition to the system performance and intelligent network services, makes the Cisco Catalyst 6500 Series unparalleled in the industry.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0502R)

C11-60010-00 01/06