

Understanding Internal MSFC Redundancy on Hybrid Mode Catalyst 6000 Switches

Document ID: 24361

Introduction

Option 1: Dual Internal MSFCs Operating as Separate Routers

Option 2: Single Router Mode

Option 3: Manual Mode Redundancy

Related Information

Introduction

This document is intended to describe the concept and role of the Designated Router (DR) with respect to internal Multilayer Switch Feature Card (MSFC) redundancy in the Catalyst 6000 platform. The configuration limitations on the internal MSFCs are discussed, along with failure scenarios of what can happen if those limitations are not followed. The advantages/disadvantages of the three types of internal MSFC redundancy options are also discussed in this document.

Option 1: Dual Internal MSFCs Operating as Separate Routers

This option was the original method of internal MSFC redundancy. When using this method, the two MSFCs operate as two separate routers. The routers must be configured within certain guidelines, and the reason for these guidelines involves the concept of the designated MSFC.

Designated MSFC

In an internally redundant MSFC configuration (a setup with two MSFCs present in the same chassis), the concept of the designated MSFC is introduced. The designated MSFC is the MSFC that comes up first, or has been up the longest. The designated MSFC can be the MSFC in slot 1 or the MSFC in slot 2. There is no mechanism to influence which MSFC will be the designated MSFC; the first to come on-line will be the designated MSFC. If the designated MSFC is reloaded manually or experiences an unexpected reload, the other MSFC will become the designated MSFC. You can verify which MSFC is the designated MSFC by issuing either the **show fm feature** or **sh redundancy** command on either MSFC.

For example, this command executed on the MSFC in slot 1 indicates that this MSFC is not the designated MSFC, and that the designated MSFC is in slot 2. The sample output is below.

```
Cat6k-MSFC-slot1#show fm feature
Redundancy Status: Non-designated
    Designated MSFC: 2
    Non-designated MSFC:1
...
```

The same command issued on the MSFC in slot 2 would display the following:

```
Cat6k-MSFC-slot2#show fm feature
Redundancy Status: designated
    Designated MSFC: 2
    Non-designated MSFC:1
```

...

The **sh redundancy** command output will display the same type of information, as shown below.

```
Cat6k-MSFC-slot1# sh redundancy
Designated Router: 2 Non-designated Router: 1
Redundancy Status: designated
...
```

Notes:

- There is no way to know in advance which MSFC will be designated.
- There is no relation between the active Supervisor (SUP) and the designated MSFC. You can have the designated MSFC in the standby SUP.
- Even in a system with a single MSFC, there will still be the concept of the designated MSFC. The designated MSFC will be the only MSFC in the chassis.
- Do not confuse the concept of the designated MSFC with the active SUP, the DR in Open Shortest Path First (OSPF), the DR in Protocol Independent Multicast (PIM), or the Hot Standby Router Protocol (HSRP) active router.

Role of the Designated MSFC

For Catalyst 6000 family switches with dual Supervisor IA (SUP IA) /Policy Feature Card (PFC)/MSFC or dual SUP IA/PFC/MSFC 2, the responsibility of the designated MSFC is as follows:

- programming the Access-List (ACL) in the hardware Ternary Content Addressable Memory (TCAM)

This induces several limitations in the MSFC configuration. The first is that both MSFCs must have the same ACL configuration and must be applied on the same VLAN interfaces. Failure to do this will result in undesired and unpredictable scenarios.

For Catalyst 6000 switches with dual SUP II/PFC 2/MSFC 2, the responsibilities of the designated MSFC are as follows:

- programming the ACL in the hardware TCAM
- downloading the Cisco Express Forwarding (CEF) table from the MSFC 2 to the hardware Forwarding Information Base (FIB) of the active PFC 2

In addition to the limitations described in the SUP IA case, there are some additional limitations. The routing table between both MSFCs needs to be the same. Failure to do that will result in unpredictable routing and switching behavior.

For example, if you have a chassis with dual Supervisor II (SUP II)/PFC 2/MSFC 2 and with the MSFC 2 in slot 1 configured correctly for routing with the expected routing table, and the MSFC 2 in slot 2 has an empty routing table. Depending on who is the designated MSFC, you may have the following behavior:

- If the MSFC 2 in slot 1 is designated, its CEF table would be downloaded to the active SUP II and the expected routing would occur.

- If the MSFC 2 in slot 2 is designated, it will not have any CEF entries, as the routing table will be empty. This will result in an empty FIB downloaded to the active SUP II and Layer 3 (L3) traffic would be dropped.

For more information about the FIB and the unicast forwarding in SUP II/PFC 2/MSFC 2 system, refer to the following:

- How-To Troubleshoot Unicast IP Routing CEF on Catalyst 6000s with a SUP 2 in Hybrid Mode

Exceptions

- ACLs are programmed only by the DR. This is valid for standard and extended security ACLs, but there are some exceptions to this rule. For example, reflexive ACLs can be programmed both by the designated MSFC and the non-designated MSFC.
- The FIB is programmed only by the DR. This is valid for all CEF entries for the network (learned per routing protocol or static routes). However, there are some exceptions as well. Some host entries such as the loopback address of a non-DR will be downloaded to the FBI per the non-DR.

Configuration Limitations

Due to the role of the designated MSFC and all the limitations described above, there are configuration restrictions on both MSFCs. Specifically, the following apply:

- Both MSFCs must have the following:
 - ◆ the same routing protocols
 - ◆ the same static routes
 - ◆ the same default routes
 - ◆ the same policy routes
 - ◆ the same VLAN interfaces
 - ◆ the same IOS ACLs applied to the same VLAN interfaces, in the same direction, on both MSFCs
 - ◆ both MSFCs should have IP addresses configured on the same subnet in the corresponding VLAN interface
- All interfaces must have the same administrative/operational status. If an interface is up on one MSFC, it must be up the second as well (cannot be shutdown on one and up on the other).

The redundancy between the two MSFCs will be provided using HSRP (usually with a different standby priority configured on each MSFC).

For L3 redundancy, the configuration of the two MSFCs should be identical, except for the following parameters:

- HSRP standby priority
- IP address commands

Advantages and Disadvantages of Option 1

Advantages

- Both MSFCs run the same routing protocols and have the same routing table. Therefore, when a failure in one MSFC occurs, the second MSFC does not need to spend time waiting for the routing protocols to converge before forwarding packets.
- HSRP can provide fast failover from active to standby in case of failure for gateway redundancy.
- Combined with high availability for Layer 2 (L2) failover, it provides recovery time within the order of few seconds in case of failure of one SUP/MSFC.

Disadvantages

- Waste of IP addresses; two IP addresses are required per VLAN and per chassis.
- Additional routing protocol peering needed.
- Non-Reverse Path Forwarding (RPF) traffic for IP multicast must be dropped in software when using the SUP IA platform.
- Complexity of maintaining two, almost identical configurations.

The last disadvantage mentioned above is addressed with the config-sync feature. Support for this feature begins with release 12.1(3a)E1 in the MSFC. For more information on config-sync, refer to MSFC Configuration Synchronization Overview.

Option 2 : Single Router Mode

Single Router Mode (SRM) is a new feature that addresses the drawback of the previous HSRP based redundancy scheme. SRM is supported starting in the following releases of the software:

- Dual SUP II/PFC 2/MSFC 2 : 12.1(8a)E2 and 6.3(1)
- Dual SUP IA/PFC/MSFC 2 : 12.1(8a)E2 and 6.3(1)
- Dual SUP IA/PFC/MSFC1 : 12.1(8a)E4 and 6.3(1)

SRM requirements:

- Both MSFCs must run the same IOS image.
- High availability needs to be configured on the SUP.
- Both MSFCs have the same configuration.
- Only the designated MSFC is seen to the network.
- The non-designated MSFC stays up with all VLAN interfaces down/down (completely booted).
- The configuration is only allowed on the designated MSFC.

When SRM is enabled, the non-DR is online, but has all of its interfaces down. Thus, it does not hold any routing table information. This means that if the DR fails, there will be some delay before the non-DR coming online will have a complete route table. To help account for this, the information being used prior to the failure by the SUP for L3 forwarding is maintained and updated with any new information from the new

DR.

SRM and SUP II/PFC 2/MSFC 2 Failure Scenario

The following will happen if the SRM and SUP II/PFC 2/MSFC 2 begin to fail:

5. The DR is failing.
4. The new DR brings up its VLAN interfaces.
3. FIB entries are maintained on the active SUP, and traffic is switched using the old FIB table for two minutes. After failure of the DR, the new DR is not allowed to update the SUP for two minutes while it is building its route table.
2. After two minutes, the new CEF table (CEF table of the new DR) is downloaded to the SUP II, whether or not the routing protocol has completed its convergence.
1. As routing protocol neighbors have their adjacencies cleared, there may still be a forwarding outage (on other devices) after the switch over.

A new feature is added in release 7.1(1) which allows the tuning of the interval between using the old FIB table and accepting the new one from the new DR. This output is displayed as follows:

```
Router(config-r-ha)#single-router-mode failover table-update-delay ?  
<0-4294967295> Delay in seconds between switch over detection and h/w FIB reload
```

Before release 7.1(1), this timer is not tunable and is always 120 seconds (two minutes). It is usually recommended to adjust the failover table-update-delay to at least the time it takes to repopulate the routing table.

SRM and SUP IA/PFC/MSFC(1 or 2) Failure Scenario

The following will happen if the SRM and SUP IA/PFC/MSFC(1 or 2) begin to fail:

4. The DR is failing.
3. The new DR brings up the VLAN interfaces
2. The existing Multilayer Switching (MLS) shortcuts are maintained on the SUP. L3 traffic continues to be routed using the old shortcut.
1. Any new flow that needs to be created is created by the new DR immediately with the following steps:
 - ◇ A packet is a candidate for the L3 shortcut.
 - ◇ The packet is forwarded to the new DR.
 - ◇ If the new DR already has a route to destination, it routes the packet and the new shortcut is created on the SUP.
 - ◇ If the new DR does not yet have a route to the destination (remember, the new DR may still be busy computing the routing table), the packet is dropped.

Advantages and Disadvantages of SRM

Advantages

- Conserves IP addresses.
- Reduces routing protocol peering.
- Configuration much simpler; no risk of running unsupported mismatched configurations

Disadvantages

- We still use the old FIB image of the routing table even though the router that creates it is not on-line anymore. There is a risk during the table-update-delay time to route packet to a non-valid route.
- Can be more disruptive to the network than Option 1, as the routing table needs to be calculated from scratch on the new DR.

Option 3: Manual Mode Redundancy

Manual Mode redundancy is no longer supported. Cisco recommends using the SRM option. Manual redundant mode involved forcing the non-designated MSFC in ROMmon mode. For more information, refer to Manual Mode MSFC Redundancy.

Related Information

- [LAN Product Support](#)
- [LAN Switching Technology Support](#)
- [Technical Support— Cisco Systems](#)

All contents are Copyright © 1992—2003 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Updated: Sep 01, 2005

Document ID: 24361
