

Cisco Catalyst 9000 Platform Graceful Insertion and Removal

Introduction

The enterprise campus network connects and enables user endpoints and all the various devices on the network. With the ongoing trends in digitization, increasing adoption of IP devices, and user applications, the network layer has become important, and so has its reliability and availability. Currently during the maintenance windows for installing modules, upgrading the software and debugging issues, there is no simpler way to isolate the device from the traffic path and hence there is an impact to the production traffic. If the device does get isolated then it requires series of commands and cumbersome process which is again not sustainable for huge ever-growing networks.

With Graceful Insertion and Removal (GIR), you can easily isolate a device from the traffic path and perform any kind of maintenance on it. GIR uses redundant paths and existing routing and First-Hop Redundancy Protocols (FHRPs) running in the network to smoothly remove a device from service for maintenance, and to return it to service when the maintenance is complete. Similar methodologies have been adopted by many organizations as part of their network maintenance best practices, but the removal and reinsertion of a device is typically performed manually. To automate this process, Cisco IOS® XE Software introduces the Graceful Insertion and Removal function on the Cisco® Catalyst® 9000 switch platform. GIR makes it easier to integrate network device maintenance into the overall network operation workflow.

This white paper provides guidelines for implementing GIR mode on Cisco Catalyst 9000 family switches for a variety of customer profiles.

Contents

Introduction

Target audience

GIR overview 2

GIR maintenance mode and templates

To move a switch into maintenance mode

To move the switch out of maintenance mode

GIR with user-defined profiles

Maintenance mode parameters

GIR snapshots and verification

GIR case studies

General topology for IS-IS and OSPF

Case study 1
(isolating IS-IS on S4)

Case study 2
(isolating OSPF on S4)

VRRP and HSRP

Case study 3
(isolating VRRP on S4)

Case study 4
(isolating HSRP on S4)

Conclusion

Additional resource

Target audience

This document is for those responsible for planning, implementation, and maintenance of networking devices.


GIR overview

GIR provides an easy way to isolate a switch for maintenance and then bring it back into service. It uses the switch's maintenance mode to systematically eject a Cisco Catalyst 9000 family switch from the network with zero or minimal disruption to network service. When a switch is in maintenance mode, it is isolated from the active forwarding paths in the network. Maintenance tasks, such as real-time debugging, hardware replacement, or software upgrade or downgrade, can be performed without affecting the production traffic. When maintenance tasks are completed, the graceful insertion function places the switch back into the network without impact.

GIR currently supports the following protocols (for both IPv4 and IPv6 address families):

- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)*
- Hot Standby Router Protocol (HSRP)*
- Virtual Router Redundancy Protocol (VRRP)*

When a Cisco Catalyst 9000 family switch is transitioning to maintenance mode for graceful removal, the routing and FHRPs running on the switch modify the routing metrics so that the neighboring devices view the maintenance switch as the least favorable candidate in their route calculation and reroute traffic to alternative paths or devices in the network.



* These protocols are supported beginning with Cisco IOS XE Release 16.8.1.

When the switch is phasing out of maintenance mode for graceful insertion, the routing and FHRP protocols reverse the changes and advertise the switch as normal. The neighboring devices perform a new route calculation and start to forward traffic to the switch if it offers the best path.

Note: GIR is intended for use in distribution or Core Layer

GIR maintenance mode and templates

A device is placed into maintenance mode after GIR isolates it from the network forwarding paths. While in maintenance mode, the device is not in network service, and hardware or software maintenance can be done on it without affecting network performance.

It takes one Cisco IOS XE privilege EXEC command, “start maintenance,” to move a Cisco Catalyst 9000 family switch into maintenance mode. This simple step allows easy operation and integration into the overall network maintenance workflow. The command triggers a predefined maintenance template, which consists of all the supported protocols and associated best practices for isolation. After all the protocols are isolated, the device enters maintenance mode.

Bringing the device out of maintenance mode requires another Cisco IOS XE privilege EXEC command, “stop maintenance.” This command moves the switch into normal mode. It triggers the same predefined maintenance template, but this time the steps are performed in reverse to return the isolated protocols to their normal states and complete the graceful insertion of the device into the network.

The following is an example of moving a switch into and out of maintenance mode using the default profile. The switch in this example is running IS-IS.

To move a switch into maintenance mode

```
Switch#start maintenance
```

```
Template default will be applied. Do you want to continue?[confirm]
```

```
Switch#
```

```
*Feb 17 02:11:49.898: %MMODE-6-MMODE_CLIENT_TRANSITION_START: Maintenance Isolate start for router isis 1
```

```
*Feb 17 02:11:49.898: %MMODE-6-MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance Isolate complete for router isis 1
```

```
*Feb 17 02:12:49.950: %MMODE-6-MMODE_ISOLATED: System is in Maintenance
```

```
Switch(maint-mode)#
```

Once the device is in maintenance mode, there is no traffic traversing through it, and the protocol neighborhood status is unchanged. Any sort of maintenance can now be performed.

To move the switch out of maintenance mode

```
Switch(maint-mode)#stop maintenance
```

```
*Feb 18 03:19:35.833: %MMODE-6-MMODE_CLIENT_TRANSITION_START: Maintenance Insert start for router isis 1
```

```
*Feb 18 03:19:35.833: %MMODE-6-MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance Insert complete for router isis 1
```

```
*Feb 18 03:20:05.974: %MMODE-6-MMODE_INSERTED: System is in Normal Mode
```

```
Switch#
```

In the default maintenance profile, forwarding protocols use different methods to isolate the switch from the active forwarding path. Table 1 shows the mechanisms used by each protocol. The goal of the mechanism is to influence the forwarding decision of the remaining devices so that they don't choose this switch as part of the best path.

Protocol	Graceful removal mechanism	Graceful insertion mechanism
IS-IS	Refresh Link-State Packets (LSPs) with overload bit	Refresh LSPs with overload bit off
OSPF	Send Link-State Advertisements (LSAs) with maximum metric	Refresh LSAs with original metric
HSRP	Advertise resign message	Advertise hello message
VRRP	Change priority to 0	Revert to original priority value

GIR with user-defined templates

In addition to the default maintenance template discussed in the previous section, Cisco IOS XE for the Cisco Catalyst 9000 family allows users to define their own templates to perform GIR on a switch. For instance, if the device has two instances of IS-IS or two protocols running on the switch, a user can define the action to isolate just one instance of IS-IS or just one protocol, and reactivate it in the normal-mode profile for graceful insertion.

The following is an example of a user-defined template. A new template is created for isolating just one IS-IS instance .

Creating a new customized template

```
switch(config)#maintenance-template User_Template
switch(config-maintenance-templ)#router isis 1
switch(config)#exit
```

switch#show system mode maintenance template User_template

```
System Mode: Normal
Maintenance Template User_Profile details:
router isis 1
```

Associate the above created template with the maintenance template

```
switch(config)#system mode maintenance
switch(config-maintenance)#template User_template
switch(config-maintenance)#exit
```

```
switch#show system mode maintenance
```

```
System Mode: Normal
```

```
Current Maintenance Parameters:
```

```
Maintenance Duration: 0 (mins)
```

```
Failsafe Timeout: 30 (mins)
```

```
Maintenance Template: User_template
```

```
Reload in Maintenance: False
```

```
switch#start maintenance
```

```
Template User_template will be applied. Do you want to continue?[confirm]
```

Once the maintenance is completed with the user-defined template, the configuration can be reverted back to the default with the sequence below:

```
switch(config)#system mode maintenance
```

```
switch(config-maintenance)#no template User_template
```

```
Core1(config-maintenance)#exit
```

As soon as the user-defined template is disassociated from the maintenance profile, it reverts back to the default template. You can either reuse the configured template later by associating it back to the maintenance profile or delete it completely from the switch by following these steps:

```
Core1(config)#no maintenance-template User_template
```

```
Core1(config)#exit
```

Maintenance mode parameters

In addition to customized templates, there are multiple configurable options in the system mode maintenance which can be changed based on certain use cases. Below is the snapshot of the options and their details:

```
switch#show system mode maintenance
```

```
System Mode: Normal
```

```
Current Maintenance Parameters:
```

```
Maintenance Duration: 0 (mins) --- The system will remain in maintenance mode till the configured time expires
```

```
Failsafe Timeout: 30 (mins) --- Configures client-ack timeout value.
```

```
If the system is going into maintenance mode, it will continue to reach maintenance. If the system is exiting from maintenance mode, then it will reach normal mode
```

```
Maintenance Template: User_template --- This shows the current template associated with the maintenance mode
```

```
Reload in Maintenance: False --- If set to true, and if the device under maintenance mode reloads, it will come back up in maintenance mode
```

GIR snapshots and verification

Cisco IOS XE uses system-generated snapshots to record the state of a switch before and after a GIR operation. A snapshot of the system status, called “before_maintenance,” is automatically generated before the switch enters GIR maintenance mode, and another snapshot, called “after_maintenance,” is automatically generated after the switch exits GIR maintenance mode—that is, returns to the normal mode. The details of the snapshots can be viewed by users or exported in XML format. The system also provides an easy way to compare the snapshots before and after the GIR maintenance operation. This offers a way to quickly evaluate the health of the switch and to verify that the switch has been inserted back to service as it should be. For example, if the traffic and route statistics are similar when you compare the snapshots taken before and after maintenance, it is a good indicator that traffic on the devices has converged properly after completion of maintenance.

The following shows sample output of the snapshot comparison.

```
Core1#show system snapshots compare before_maintenance after_maintenance
```

```
=====
```

Feature	Tag	.before_maintenance	.after_maintenance
=====			
[interface]			

[Name:GigabitEthernet0/0]			
Lastinput		00:00:01,	**00:00:24,**
Drops		1178	**1179**
packetsinput		236242	**236378**
[Name:FortyGigabitEthernet1/0/7]			
output		00:00:14,	**00:00:48,**
packetsinput		97457685	**97457722**
packetsoutput		97305093	**97305102**
[Name:FortyGigabitEthernet1/0/8]			
output		00:00:03,	**00:00:44,**
packetsinput		97464520	**97464557**
packetsoutput		97297816	**97297824**
[route-summary]			

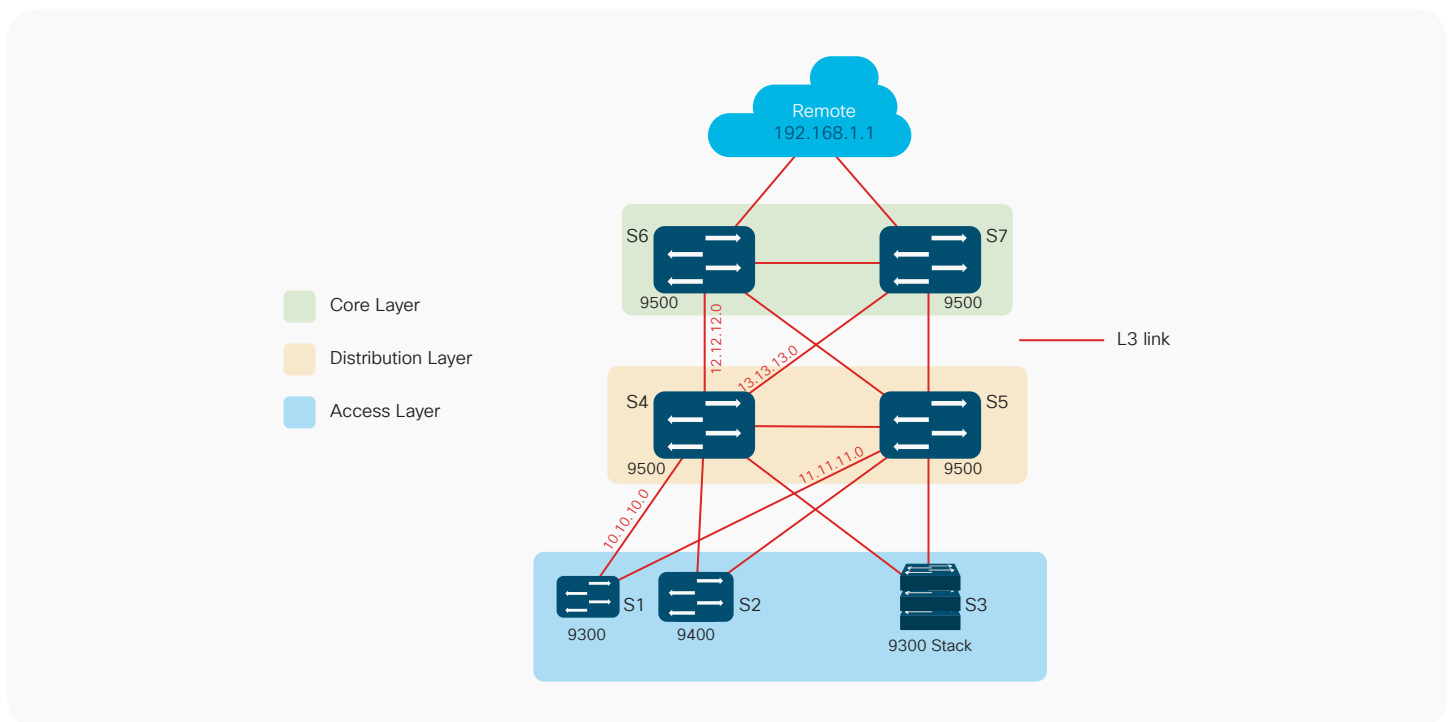
[routev6-summary]			

GIR case studies

This section demonstrates GIR operation for different network protocols using the system default profiles.

General topology for campus running routing protocols

Figure 1 shows the topology used for the IS-IS and OSPF protocols.



Case study 1 (isolating IS-IS on distribution Switch S4)

IS-IS graceful removal is done by setting the overload bit in the Link State Packets (LSPs). When a switch sees this bit set in the LSP from its peer, the switch does not send traffic toward that peer. IS-IS graceful insertion is done by unsetting the overload bit in the LSP.

In this example, IS-IS is configured between all Layer 3 nodes. The S1 switch is learning the remote prefix 192.168.1.1 from both the S4 and S5 switches. After graceful removal has been performed on S4 with the “start maintenance” command, all remote subnets are advertised with a maximum metric of 83; however, the subnets connected to S4 will have the original metric. This will force S1 to remove the external route to 192.168.1.1 that was pointing to S4 from its routing table.

S4 in normal mode

```
S1#show ip route 192.168.1.1
Routing entry for 192.168.1.1/32
  Known via "isis", distance 115, metric 40, type level-1
  Redistributing via isis
  Last update from 11.11.11.2 on
  FortyGigabitEthernet1/1/2, 00:00:15 ago
  Routing Descriptor Blocks:
    11.11.11.2, from 192.168.1.1, 00:00:15
    ago, via FortyGigabitEthernet1/1/2
      Route metric is 40, traffic share
      count is 1
    * 10.10.10.2, from 192.168.1.1, 00:00:15
    ago, via FortyGigabitEthernet1/1/1
      Route metric is 40, traffic share
      count is 1

S4#show ip route 192.168.1.1
Routing entry for 192.168.1.1/32
  Known via "isis", distance 115,
  metric 30, type level-1
  Redistributing via isis
  Last update from 13.13.13.2 on
  FortyGigabitEthernet1/0/2, 00:01:05 ago
  Routing Descriptor Blocks:
    13.13.13.2, from 192.168.1.1, 00:01:05
    ago, via FortyGigabitEthernet1/0/2
      Route metric is 30, traffic share
      count is 1
    * 12.12.12.2, from 192.168.1.1, 00:01:05
    ago, via FortyGigabitEthernet1/0/1
      Route metric is 30, traffic share
      count is 1
```

S4 in maintenance mode

```
S1#show ip route 192.168.1.1
Routing entry for 192.168.1.1/32
  Known via "isis", distance 115, metric 40, type level-1
  Redistributing via isis
  Last update from 11.11.11.2 on
  FortyGigabitEthernet1/1/2, 00:00:09 ago
  Routing Descriptor Blocks:
    * 11.11.11.2, from 192.168.1.1, 00:00:09
    ago, via FortyGigabitEthernet1/1/2
      Route metric is 40, traffic share
      count is 1

S4(maint-mode)#show ip route 192.168.1.1
Routing entry for 192.168.1.1/32
  Known via "isis", distance 115, metric
  83, type level-1
  Redistributing via isis
  Last update from 13.13.13.2 on
  FortyGigabitEthernet1/0/2, 00:01:09 ago
  Routing Descriptor Blocks:
    13.13.13.2, from 192.168.1.1, 00:01:09
    ago, via FortyGigabitEthernet1/0/2
      Route metric is 83, traffic share
      count is 1
    * 12.12.12.2, from 192.168.1.1, 00:01:09
    ago, via FortyGigabitEthernet1/0/1
      Route metric is 83, traffic share
      count is 1
```


S4 in normal mode	S4 in maintenance mode
<pre> S4#show isis rib 192.168.1.1 IPv4 local RIB for IS-IS process IPv4 unicast topology base (TID 0, TOPOID 0x0) ===== Routes under majornet 192.168.1.0/24: 192.168.1.1/32 prefix attr X:0 R:0 N:0 [115/L1/30] via 12.12.12.2(FortyGigabitEthernet1/0/1), from 192.168.1.1, tag 0, LSP[12/18/6] prefix attr: X:0 R:0 N:0 (installed) [115/L1/30] via 13.13.13.2(FortyGigabitEthernet1/0/2), from 192.168.1.1, tag 0, LSP[21/18/6] prefix attr: X:0 R:0 N:0 (installed) </pre>	<pre> S4(maint-mode)#show isis rib 192.168.1.1 IPv4 local RIB for IS-IS process IPv4 unicast topology base (TID 0, TOPOID 0x0) ===== Routes under majornet 192.168.1.0/24: 192.168.1.1/32 prefix attr X:0 R:0 N:0 [115/L1/83] via 12.12.12.2(FortyGigabitEthernet1/0/1), from 192.168.1.1, tag 0, LSP[12/18/7] prefix attr: X:0 R:0 N:0 (installed) [115/L1/83] via 13.13.13.2(FortyGigabitEthernet1/0/2), from 192.168.1.1, tag 0, LSP[21/18/7] prefix attr: X:0 R:0 N:0 (installed) </pre>

Case study 2 (isolating OSPF on distribution Switch S4)

OSPF graceful removal is done by refreshing the Link-State Advertisement (LSA) with metric 65537. When other switches see this LSA, they will no longer send traffic toward this device. OSPF graceful insertion is done by refreshing the LSA with the original metric. In this example, OSPF is configured between all Layer 3 nodes. The S1 switch is learning the remote prefix 192.168.1.1 from both the S4 and S5 switches.

After graceful removal has been performed on S4 with the “start maintenance” command, all remote subnets are advertised with a maximum metric of 65537; however, the subnets connected to S4 will have the original metric. This will force S1 to remove the external route to 192.168.1.1 that was pointing to S4 from its routing table.

S4 in normal mode

```
S1#show ip route 192.168.1.1
Routing entry for 192.168.1.0/24
  Known via "ospf 1", distance 110,
  metric 4, type intra area
  Last update from 10.10.10.2 on
  FortyGigabitEthernet1/1/1, 01:31:48 ago
Routing Descriptor Blocks:
  * 11.11.11.2, from 192.168.1.1, 2d17h
  ago, via FortyGigabitEthernet1/1/2
    Route metric is 4, traffic share
    count is 1
  10.10.10.2, from 192.168.1.1, 01:31:48
  ago, via FortyGigabitEthernet1/1/1
    Route metric is 4, traffic share
    count is 1
S4#show ip route 192.168.1.1
Routing entry for 192.168.1.0/24
  Known via "ospf 1", distance 110,
  metric 3, type intra area
  Last update from 13.13.13.2 on
  FortyGigabitEthernet1/0/2, 01:34:55 ago
Routing Descriptor Blocks:
  13.13.13.2, from 192.168.1.1, 01:34:55
  ago, via FortyGigabitEthernet1/0/2
    Route metric is 3, traffic share
    count is 1
  * 12.12.12.2, from 192.168.1.1, 01:34:55
  ago, via FortyGigabitEthernet1/0/1
    Route metric is 3, traffic share
    count is 1
```

S4 in maintenance mode

```
S1#show ip route 192.168.1.1
Routing entry for 192.168.1.0/24
  Known via "ospf 1", distance 110, metric
  4, type intra area
  Last update from 11.11.11.2 on
  FortyGigabitEthernet1/1/2, 01:43:27 ago
Routing Descriptor Blocks:
  * 11.11.11.2, from 192.168.1.1, 2d17h
  ago, via FortyGigabitEthernet1/1/2
    Route metric is 4, traffic share
    count is 1
S4(maint-mode)#show ip route 192.168.1.1
Routing entry for 192.168.1.0/24
  Known via "ospf 1", distance 110, metric
  65537, type intra area
  Last update from 13.13.13.2 on
  FortyGigabitEthernet1/0/2, 00:07:45 ago
Routing Descriptor Blocks:
  13.13.13.2, from 192.168.1.1, 00:07:45
  ago, via FortyGigabitEthernet1/0/2
    Route metric is 65537, traffic share
    count is 1
  * 12.12.12.2, from 192.168.1.1, 00:07:45
  ago, via FortyGigabitEthernet1/0/1
    Route metric is 65537, traffic share
    count is 1
```

S4 in normal mode

```
S4#show ip ospf rib 192.168.1.1
      OSPF Router with ID
(13.13.13.1) (Process ID 1)
      Base Topology (MTID 0)

OSPF local RIB

Codes: * - Best, > - Installed in global
RIB

LSA: type/LSID/originator

*> 192.168.1.0/24, Intra, cost 3, area 0
      SPF Instance 17, age 01:36:34
      Flags: RIB
      via 12.12.12.2,
FortyGigabitEthernet1/0/1, label 1048578,
strict label 1048578
      Flags: RIB
      LSA: 1/192.168.1.1/192.168.1.1
      via 13.13.13.2,
FortyGigabitEthernet1/0/2, label 1048578,
strict label 1048578
      Flags: RIB
      LSA: 1/192.168.1.1/192.168.1.1
```

S4 in maintenance mode

```
S4(maint-mode)#show ip ospf rib
192.168.1.1
      OSPF Router with ID
(13.13.13.1) (Process ID 1)
      Base Topology (MTID 0)

OSPF local RIB

Codes: * - Best, > - Installed in global
RIB

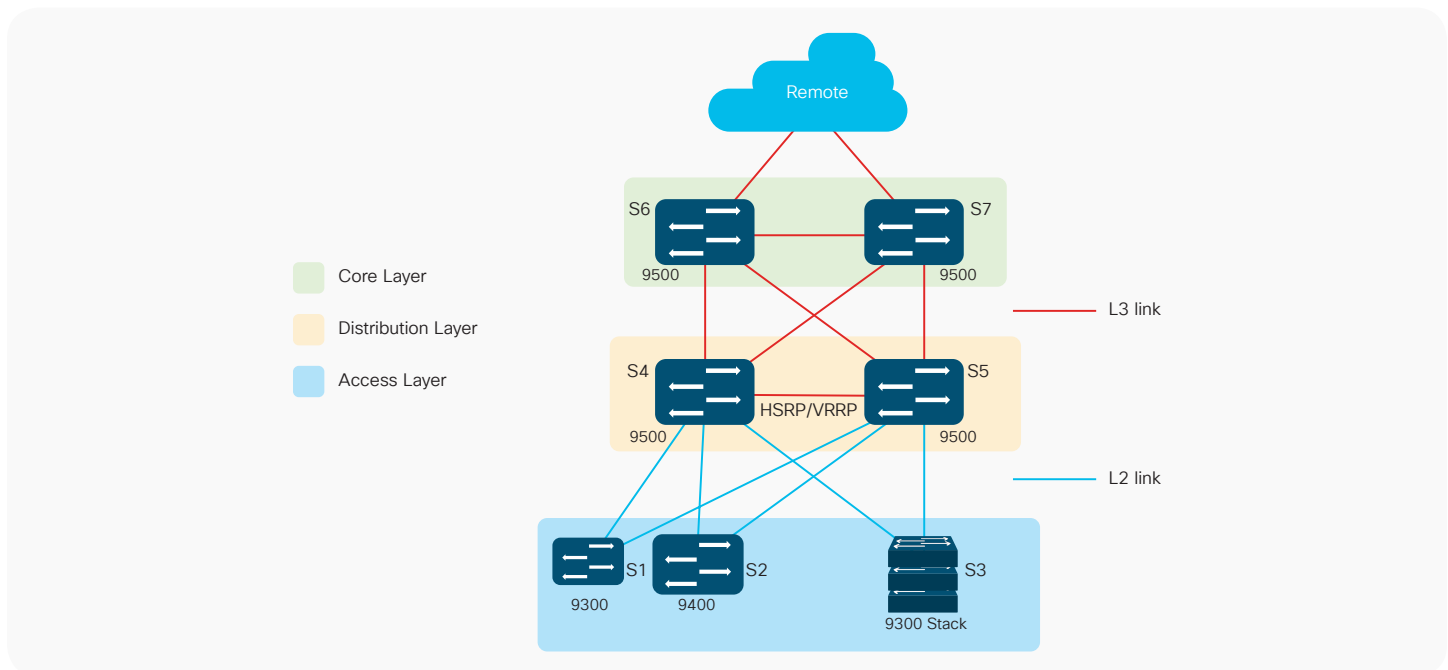
LSA: type/LSID/originator

*> 192.168.1.0/24, Intra, cost 65537,
area 0
      SPF Instance 18, age 00:08:39
      Flags: RIB
      via 12.12.12.2,
FortyGigabitEthernet1/0/1, label 1048578,
strict label 1048578
      Flags: RIB
      LSA: 1/192.168.1.1/192.168.1.1
      via 13.13.13.2,
FortyGigabitEthernet1/0/2, label 1048578,
strict label 1048578
      Flags: RIB
      LSA: 1/192.168.1.1/192.168.1.1
```

VRRP and HSRP

Below topology represents a typical multi-tier design comprising of Layer 2 connections between Access and Distribution and Layer 3 connections between Distribution and Core. Redundancy is achieved by implementing protocols like HSRP or VRRP in the distribution layer.

Figure 2. GIR Topology for VRRP and HSRP



Case study 3 (isolating VRRP on distribution Switch S4)

VRRP graceful shutdown is done by setting the priority value to zero, which forces the backup switch to become the master immediately. The VRRP master device goes into the INIT state and stops advertising hello messages completely. During this process, there is very little impact to the end user traffic. If the interfaces between S4/S5 and S1/S2/S3 are Layer 2 enabled, a customized template has to be created to exclude the Layer 2 shutdown; if the interfaces are Layer 3, maintenance mode can be executed with the default template.

In this example, VRRP is configured between S4 and S5, with S4 being the master and the S5 being the backup.

After graceful removal has been performed on S4 with the “start maintenance” command, S5 becomes the master and starts servicing all the endpoints coming from the access layer devices.

S4 in normal mode

```
S4#show system mode maintenance template  
vrrptest
```

Maintenance Template vrrptest details:

```
vrrpv3 Vlan100 1 ipv4
```

```
S4#show vrrp brief
```

Interface	Grp	A-F	Pri	Time	
Own	Pre	State	Master	addr/Group	addr
Vl100	1	IPv4	110	0	
N	Y	MASTER	100.100.100.2(local)		
			100.100.100.1		

```
S4#show vrrp detail
```

Vlan100 - Group 1 - Address-Family IPv4

State is MASTER

State duration 38.989 secs

Virtual IP address is 100.100.100.1

Virtual MAC address is 0000.5E00.0101

Advertisement interval is 1000 msec

Preemption enabled

Priority is 110

**Master Router is 100.100.100.2 (local),
priority is 110**

Master Advertisement interval is 1000
msec (expires in 321 msec)

S4 in maintenance mode

```
S4(maint-mode)#show vrrp brief
```

Interface	Grp	A-F	Pri	Time	
Own	Pre	State	Master	addr/Group	addr
Vl100	1	IPv4	110	0	N
Y	INIT		100.100.100.3	100.100.100.1	

```
S4(maint-mode)#show vrrp detail
```

Vlan100 - Group 1 - Address-Family IPv4

**State is INIT (Group is shutdown - In
maintenance mode)**

State duration 49.234 secs

Virtual IP address is 100.100.100.1

Virtual MAC address is 0000.5E00.0101

Advertisement interval is 1000 msec

Preemption enabled

Priority is 110

**Master Router is 100.100.100.3, priority
is 100**

Master Advertisement interval is 1000
msec (learned)

Master Down interval is unknown

VRRPv3 Advertisements: sent 81135
(errors 0) - rcvd 1762

S4 in normal mode

```
Master Down interval is unknown
  VRRPv3 Advertisements: sent 80951
(errors 0) - rcvd 1762
  VRRPv2 Advertisements: sent 0 (errors 0)
- rcvd 0
  Group Discarded Packets: 0
  VRRPv2 incompatibility: 0
  IP Address Owner conflicts: 0
  Invalid address count: 0
  IP address configuration mismatch : 0
  Invalid Advert Interval: 0
  Adverts received in Init state: 0
  Invalid group other reason: 0
  Group State transition:
    Init to master: 0
    Init to backup: 8 (Last change Wed Feb
28 04:26:01.770)
    Backup to master: 9 (Last change Wed
Feb 28 04:26:05.340)
    Master to backup: 1 (Last change Tue
Feb 27 06:32:17.217)
    Master to init: 7 (Last change Wed Feb
28 03:44:00.473)
  Backup to init: 0
```

S4 in maintenance mode

```
VRRPv2 Advertisements: sent 0 (errors
0) - rcvd 0
  Group Discarded Packets: 0
  VRRPv2 incompatibility: 0
  IP Address Owner conflicts: 0
  Invalid address count: 0
  IP address configuration mismatch : 0
  Invalid Advert Interval: 0
  Adverts received in Init state: 0
  Invalid group other reason: 0
  Group State transition:
    Init to master: 0
    Init to backup: 8 (Last change Wed Feb
28 04:26:01.770)
    Backup to master: 9 (Last change Wed
Feb 28 04:26:05.340)
    Master to backup: 1 (Last change Tue
Feb 27 06:32:17.217)
    Master to init: 8 (Last change Wed Feb
28 04:29:27.898)
```

Case study 4 (isolating HSRP on Distribution Switch S4)

HSRP graceful shutdown is done by advertising a resign message on the active device, which forces the standby switch to become active immediately. The HSRP active device goes into the INIT state and stops advertising hello messages to disqualify itself from the election. During this process, there is very little impact to end user traffic. If the interfaces between S4/S5 and S1/S2/S3 are Layer 2 enabled, a customized template must be created to exclude the Layer 2 shutdown; if the interfaces are Layer 3, maintenance mode can be executed with the default template.

In this example, VRRP is configured between S4 and S5, with S4 being the active and the S5 being the standby.

After graceful removal has been performed on S4 with the “start maintenance” command, S5 becomes the master and starts servicing all the endpoints coming from the access layer devices.

S4 in normal mode	S4 in maintenance mode
<pre>S4#show system mode maintenance template hsrptest Maintenance Template hsrptest details: hsrp Vlan100 1 Master to init: 7 (Last change Wed Feb 28 03:44:00.473) Backup to init: 0 S4#show standby all Vlan100 - Group 1 State is Active 4 state changes, last state change 00:00:40 Virtual IP address is 100.100.100.1 Active virtual MAC address is 0000.0c07. ac01 (MAC In Use) Local virtual MAC address is 0000.0c07.ac01 (v1 default) Hello time 3 sec, hold time 10 sec Next hello sent in 0.432 secs Preemption enabled Active router is local Standby router is 100.100.100.3, priority 100 (expires in 8.800 sec) Priority 110 (configured 110) Group name is "hsrp-Vl100-1" (default)</pre>	<pre>S4(maint-mode)#show standby all Vlan100 - Group 1 State is Init (Maintenance mode) 5 state changes, last state change 00:04:38 Virtual IP address is 100.100.100.1 Active virtual MAC address is 0000.0c07. ac01 (MAC Not In Use) Local virtual MAC address is 0000.0c07.ac01 (v1 default) Hello time 3 sec, hold time 10 sec Preemption enabled Active router is 100.100.100.3, priority 100 (expires in 9.824 sec) Standby router is unknown Priority 110 (configured 110) Group name is "hsrp-Vl100-1" (default)</pre>

Conclusion

The GIR feature lets the network operator define an operational mode for Cisco Catalyst 9000 devices that allows the operator to perform graceful removal and insertion of devices with little to zero service disruption. When a device is in GIR mode and is not actively forwarding any traffic, the network operator can perform the majority of maintenance tasks, including software upgrades and reloads. Without the GIR feature, network operators need to perform many steps and enter many commands manually to isolate and then return a device to the network, which is a cumbersome and error-prone process.

Additional resource

High Availability Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/16-6/configuration_guide/b_166_ha_9500/b_166_ha_9500_chapter_010.html