# Cisco Expert-Level Training for CCIE Service Provider Exercise Workbook Fast Convergence Lab 01 Scenario and Answer Key

The Cisco Expert-Level Training for CCIE Service Provider Workbook contains challenging scenarios at the CCIE level that can be used for rigorous self-paced practice.

 

# Cisco 360 CCIE Training for CCIE Service Provider Exercise Workbook Fast Covergence Lab 01 Scenario and Answer Key

# Table of Contents



 

# Activity Objectives

This exercise is intended to demonstrate some ways to configure a network, as well as to help you identify if you are prepared to face the CCIE Service Provider lab exam.

Bear in mind, you will not see every feature and topic described in the exam blueprint. Furthermore, not all topics seen in lab sections appear on every exam.

Please, ensure that you consult all current official Cisco documentation before proceeding with a design or installation. This lab is primarily intended to be a learning tool, and may not necessarily follow best practice recommendation at all times, in order to convey specific information.

# General Lab Instructions

Read the following instructions carefully to ensure that you will have the lab working as expected:

- Get access to the lab and create the topology and paste the STARTUP configuration that is provided with this workbook. The access to the lab is not provided with this workbook. You can rent physical equipment or use a simulator environment that supports Cisco IOS, Cisco IOS-XE, and Cisco IOS-XR images.
- Depending on which devices you are using, the interface name and number might be different, ensure that you update the interface name and numbering, before you paste the STARTUP configuration provided.
- In the lab, you should be able to access all devices via Telnet.
- Begin your lab practice, checking the following for each device:
  - Hostname
  - Prompt login for Cisco IOS-XR only.
    - username and password for login is “cisco”.
  - For Cisco IOS and Cisco IOS-XE devices no password is required to access the privilege level 15
  - Verify that all devices see the neighbor devices according to the topology provided. Use **show cdp neighbor** command for this verification.

This lab is split into 10 different fast convergence solutions.

It is recommended to use **debug** commands and/or **show** commands as much as you can. The use of **show run** command output will not always help you to speed up identifying what is applied, especially when the whole configuration involves many pages. Unless you use the outputs with filters to help you verify specific information in the configuration section.

The following configuration are already applied on startup config:

- IPv4 addresses for all Physical and Loopback interfaces
- IS-IS level 2-only enabled in all devices
- MPLS LDP enabled in all devices
- BGP configured.
  - EBGP IPv4 session: between RT7 and RT1-XR, and between RT5 and RT4
  - IBGP RR: RT2 and RT3
  - IBGP route-reflector client: RT1-XR, RT4, and RT5.

 

# Topology

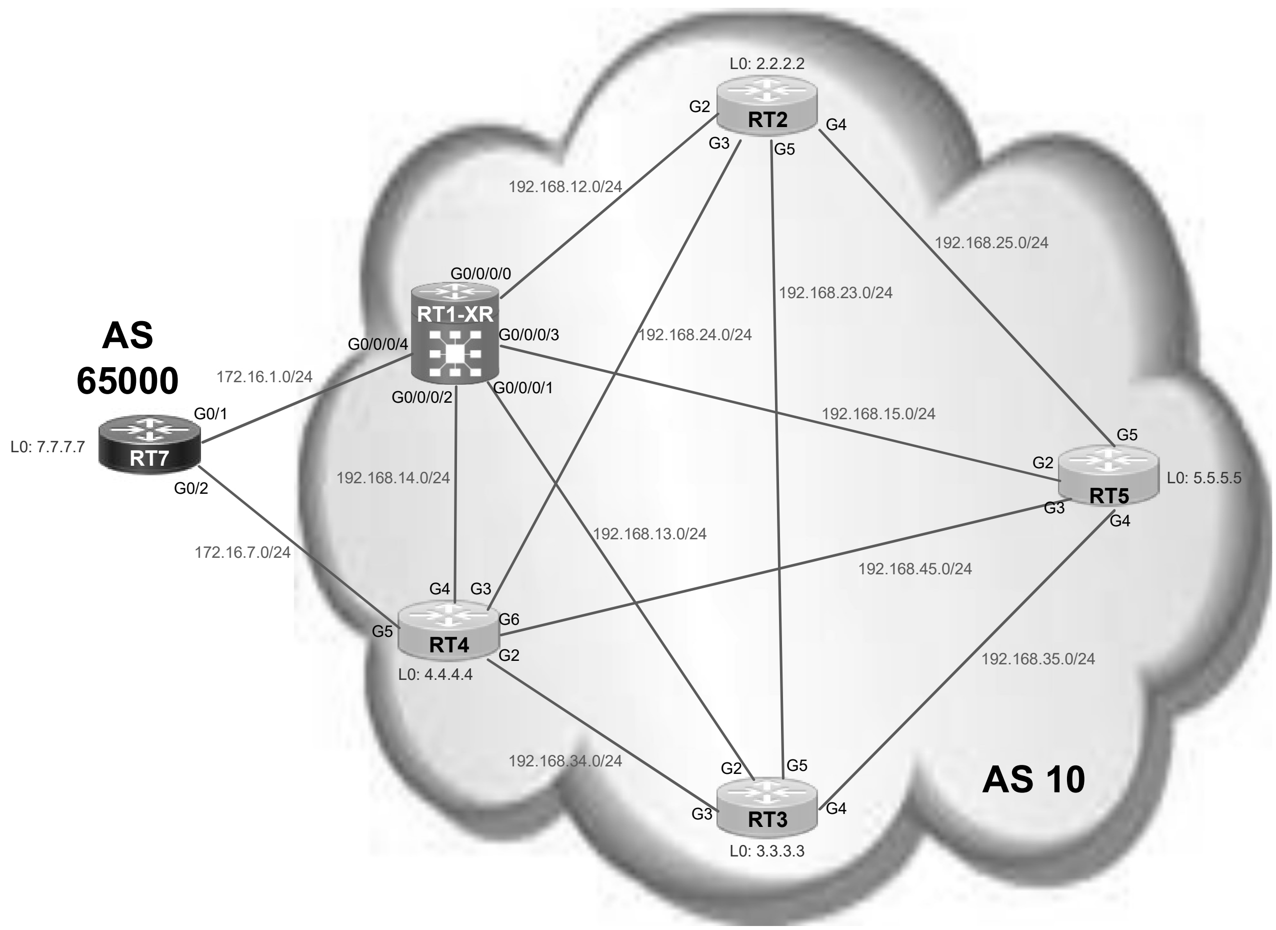


**Figure 1 - Lab Topology and IPv4 addresses**

 

# Lab Instructions

## 1. BFD

**QUESTION:**

It is taking too long for the layer 3 to realize that there was a signal loss on the transport layer on the link between RT2 and RT4. Consequently, this is impacting application performance during IGP convergence. Implement a fast convergence mechanism so that the link failure can be detected within 500ms.

**SOLUTION:**

The fast link failure detection can be achieved through the BDF protocol. The interval and multiplier values can be different between the routers, however hold timer value cannot be higher than 500ms (as stated in the question).

**CONFIGURATION:**

In order to enable BFD, you must enable BFD under the interface configuration mode and also under the IGP process configuration mode.

On RT2 configure the following:

```
bfd-template single-hop RT2-RT4
  interval min-tx 150 min-rx 150 multiplier 3
  !
interface GigabitEthernet 3
  bfd template RT2-RT4
!
router isis lab
  bfd all-interfaces
```

On RT4 configure the following:

```
bfd-template single-hop RT2-RT4
  interval min-tx 150 min-rx 150 multiplier 3
  !
interface GigabitEthernet 3
  bfd template RT2-RT4
!
router isis lab
  bfd all-interfaces
```

The **bfd-template** command was used as an alternative to apply BFD configuration under the interface. This is very useful when you need to enable BFD in many interfaces and you would like to ensure that you will be using the same parameters to all interfaces.

☞ **Note:** BFD template by default uses non-echo mode.

 

## VALIDATION:

Check if the interface has established BFD peering.

```
RT4# show bfd neighbors
IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
192.168.24.2                            1/1          Up        Up        Gi3
```

```
RT2# show bfd neighbors

IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
192.168.24.4                            1/1          Up        Up        Gi3
```

This is another way to verify if BFD is enabled correctly.

```
RT2# show isis neighbors detail
Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi2           192.168.12.1    UP    24       00
  Area Address(es): 49.1234
  SNPA: fa16.3e58.15e3
  State Changed: 00:36:12
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
  BFD enabled: (MTID:0, ipv4)
RT4             L2   Gi3           192.168.24.4    UP    25       02
  Area Address(es): 49.1234
  SNPA: fa16.3e65.4cab
  State Changed: 00:01:12
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  Remote BFD Support:TLV (MTID:0, IPV4)
  BFD enabled: (MTID:0, ipv4)
--snip--
```

```
RT2# show clns interface
GigabitEthernet1 is administratively down, line protocol is down
  CLNS protocol processing disabled
GigabitEthernet2 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 11 seconds
  Routing Protocol: IS-IS (lab)
    Circuit Type: level-1-2
    Interface number 0x1, local circuit ID 0x101
    Neighbor System-ID: RT1-XR
    Level-2 Metric: 25, Priority: 64, Circuit ID: RT2.01
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 7 seconds
    if state UP
  BFD enabled: (MTID:0, ipv4)
GigabitEthernet3 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 36 seconds
```

 

```
  Routing Protocol: IS-IS (lab)
    Circuit Type: level-1-2
    Interface number 0x2, local circuit ID 0x102
    Neighbor System-ID: RT4
    Level-2 Metric: 10, Priority: 64, Circuit ID: RT4.02
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 5 seconds
    if state UP
  BFD enabled: (MTID:0, ipv4)
--snip--
Loopback0 is up, line protocol is up
  Checksums enabled, MTU 1514, Encapsulation LOOPBACK
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching disabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 29 seconds
  Routing Protocol: IS-IS (lab)
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x100
    Level-2 Metric: 10, Priority: 64, Circuit ID: RT2.00
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 0
    Next IS-IS Hello in 0 seconds
    if state DOWN
  BFD enabled: (MTID:0, ipv4)
```

```
RT4# show bfd neighbors details
IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
192.168.24.2                            1/1          Up        Up        Gi3
Session state is UP and not using echo function.
Session Host: Hardware
OurAddr: 192.168.24.4
Handle: 1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 150000, MinRxInt: 150000, Multiplier: 3
Received MinRxInt: 150000, Received Multiplier: 3
Holddown (hits): 0(0), Hello (hits): 150(0)
Rx Count: 15234, Rx Interval (ms) min/max/avg: 72/514/142
Tx Count: 15218, Tx Interval (ms) min/max/avg: 123/621/142
Elapsed time watermarks: 0 0 (last: 0)
Registered protocols: ISIS CEF
Template: RT2-RT4
Uptime: 00:35:49
Last packet: Version: 1                  - Diagnostic: 0
             State bit: Up               - Demand bit: 0
             Poll bit: 0                 - Final bit: 0
             C bit: 1
             Multiplier: 3               - Length: 24
             My Discr.: 1                - Your Discr.: 1
             Min tx interval: 150000     - Min rx interval: 150000
             Min Echo interval: 0
```

### TROUBLESHOOTING:

☞ Note that, if you forget to apply the **bfd all-interfaces** command under the IS-IS process configuration mode, the BFD neighborship will not establish, even though BFD is configured under the interface. You must enable it under the client protocol, which in this case is IS-IS, and you must also configure it under the interface configuration mode, which defines the timers' parameters.

 

In this example, we removed the **bfd all-interfaces** command under the IS-IS process on RT2. Now, we will verify which impact this causes. Use this methodology in removing one command to understand the messages of the output. In doing this, you can learn more how to troubleshoot the technology and identify quickly what it is missing in the configuration or what is wrongly configured.

As you can see, now, there is no BFD peering between RT2 and RT4.

```
RT2# show bfd neighbors

RT2#
```

☞ Note that, there isn't any information with regards to BFD on the output of the **show isis neighbor details** command, except the information learnt from the neighbor. Compare this output to the previous output when everything was configured correctly and notice the differences.

```
RT2# show isis neighbors detail
Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi2           192.168.12.1    UP    23       00
  Area Address(es): 49.1234
  SNPA: fa16.3e58.15e3
  State Changed: 00:51:46
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
RT4             L2   Gi3           192.168.24.4    UP    21       02
  Area Address(es): 49.1234
  SNPA: fa16.3e65.4cab
  State Changed: 00:16:45
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  Remote BFD Support:TLV (MTID:0, IPV4)
```

However, on RT4 side, you can see the information about BFD for the **show bfd neighbors details** command, even though there isn't any BFD peering. This indicates that there is something missing on RT2 side. RT4 side looks almost ok, just missing the info "Remote BFD Support" entry with regards to RT2 neighbor.

```
RT4# show bfd neighbors

RT2#
```

```
RT4# show isis neighbors detail

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi4           192.168.14.1    UP    22       00
  Area Address(es): 49.1234
  SNPA: fa16.3e7c.4725
  State Changed: 00:53:21
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet4
  BFD enabled: (MTID:0, ipv4)
RT2             L2   Gi3           192.168.24.2    UP    26       02
  Area Address(es): 49.1234
  SNPA: fa16.3e99.242e
  State Changed: 00:18:20
  Format: Phase V
```

 

```
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  BFD enabled: (MTID:0, ipv4)
RT3             L2   Gi2            192.168.34.3    UP    26        02
  Area Address(es): 49.1234
  SNPA: fa16.3e7d.ad59
  State Changed: 01:05:12
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
  BFD enabled: (MTID:0, ipv4)
RT5             L2   Gi6            192.168.45.5    UP    23        02
  Area Address(es): 49.1234
  SNPA: fa16.3e26.c3d0
  State Changed: 01:05:12
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet6
  BFD enabled: (MTID:0, ipv4)
```

☞ Note that, on the router that does not have any BFD configuration, it shows something, we used RT3 as example to verify this.

```
RT3# show isis neighbors detail

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi2           192.168.13.1    UP    23       00
  Area Address(es): 49.1234
  SNPA: fa16.3e26.7c8b
  State Changed: 01:07:46
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
RT4             L2   Gi3           192.168.34.4    UP    28       01
  Area Address(es): 49.1234
  SNPA: fa16.3e4c.6b73
  State Changed: 00:09:53
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  Adjacency Not Up: Waiting For BFD Session
  Remote BFD Support:TLV (MTID:0, IPV4)
RT5             L2   Gi4           192.168.35.5    UP    29       03
  Area Address(es): 49.1234
  SNPA: fa16.3ee9.8ff5
  State Changed: 01:19:37
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet4
```

Conclusion, BFD is enabled under the client protocol , which in this case is IS-IS. The information applied under the interface is a customization of how BFD will operate. However, you must still configure it, as well.

Just to see the results, we applied the **bfd all-interfaces** command under the IS-IS process for RT3, however we didn't apply any BFD command under the interface. Let's see the impact.

```
RT3# show bfd neighbors

RT3#
```

```
RT3# show isis neighbors detail

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi2           192.168.13.1    UP    21       00
  Area Address(es): 49.1234
  SNPA: fa16.3e26.7c8b
  State Changed: 01:12:57
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
  BFD enabled: (MTID:0, ipv4)
RT4             L2   Gi3           192.168.34.4    DOWN  24       01
  Area Address(es): 49.1234
  SNPA: fa16.3e4c.6b73
  State Changed: 00:00:22
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  Adjacency Not Up: Waiting For BFD Session
  Remote BFD Support:TLV (MTID:0, IPV4)
  BFD enabled: (MTID:0, ipv4)
RT5             L2   Gi4           192.168.35.5    UP    24       03
  Area Address(es): 49.1234
  SNPA: fa16.3ee9.8ff5
  State Changed: 01:24:48
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet4
  BFD enabled: (MTID:0, ipv4)
```

Here is how it looks on RT4 side:

```
RT4# show isis neighbors detail

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi4           192.168.14.1    UP    25       00
  Area Address(es): 49.1234
  SNPA: fa16.3e7c.4725
  State Changed: 01:14:11
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet4
  BFD enabled: (MTID:0, ipv4)
RT2             L2   Gi3           192.168.24.2    UP    25       02
  Area Address(es): 49.1234
  SNPA: fa16.3e99.242e
  State Changed: 00:39:11
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet3
  Remote BFD Support:TLV (MTID:0, IPV4)
  BFD enabled: (MTID:0, ipv4)
RT3             L2   Gi2           192.168.34.3    DOWN  24       02
  Area Address(es): 49.1234
  SNPA: fa16.3e7d.ad59
  State Changed: 00:01:36
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet2
  Adjacency Not Up: Waiting For BFD Session
```

 

```
  Remote BFD Support:TLV (MTID:0, IPV4)
  BFD enabled: (MTID:0, ipv4)
RT5             L2   Gi6           192.168.45.5    UP    23        02
  Area Address(es): 49.1234
  SNPA: fa16.3e26.c3d0
  State Changed: 01:26:02
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet6
  BFD enabled: (MTID:0, ipv4)
```

Most interesting point is, if you apply the BDF command under the ISIS process and you forget to apply the BFD parameters under the interface, the IS-IS neighborship status will be down.

```
RT4# show isis neighbor

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi4           192.168.14.1    UP    21        00
RT2             L2   Gi3           192.168.24.2    UP    26        02
RT3             L2   Gi2           192.168.34.3    DOWN  23        02
RT5             L2   Gi6           192.168.45.5    UP    23        02
```

```
RT2# show isis neighbors

Tag lab:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
RT1-XR          L2   Gi2           192.168.13.1    UP    23        00
RT4             L2   Gi3           192.168.34.4    DOWN  20        01
RT5             L2   Gi4           192.168.35.5    UP    22        03
```

☞ **Note**: BFD does not work on an Cisco IOS XRv image that runs on a virtualization platform. This is because BFD on Cisco IOS XR is implemented on line cards (hardware dependent feature). All hardware dependent feature at the moment are not supported on any virtualization platform.

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-2/interfaces/configuration/guide/hc42xr12kbook/hc42bifw.html

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/15-mt/irb-15-mt-book/irb-bi-fwd-det.html

# 2. IP Fast Reroute (FRR) Loop-Free Alternate (LFA)

**QUESTION:**

An operation engineer noticed that during a link or node failure some routers in the core network take a long time to converge due to high CPU utilization. Enhance the IS-IS routing protocol on all devices for AS 10 (RT1-XR, RT2, RT3, RT4, and RT5), so that the backup path can be pre-computed (If there is any).

**SOLUTION:**

The IS-IS LFA feature meets this requirement by pre-computing a backup path and storing in the CEF table. It is normal that not all prefixes will have a backup path, it depends on the LFA algorithm. To know more about how LFA works, read the reference listed at the end of this section.

**CONFIGURATION:**

In order to enable IP FRR, you just need to enable it under the IGP routing process.

On RT2, RT3, RT4, and RT5 configure the following:

```
router isis lab
  fast-reroute per-prefix level-2 all
```

On RT1-XR configure the following:

```
group IPFRR
 router isis '.*'
  interface 'Giga.*'
   address-family ipv4 unicast
    fast-reroute per-prefix level 2
!
router isis lab
 apply-group IPFRR
```

The **group** command was used as an alternative to apply IP FRR under the IS-IS process enabling it for all GigabitEthernet interfaces used on the IS-IS process. This minimizes repetitive configuration using a regular expression. More information about how to use the group command, see the reference listed at the end of this section.

**VALIDATION:**

Checking the Cisco IOS XR configuration that used the group command.

```
RP/0/0/CPU0:RT1-XR# show run router isis
router isis lab
 apply-group IPFRR
 is-type level-2-only
 net 49.1234.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
```

```
  !
 !
 interface GigabitEthernet0/0/0/0
  point-to-point
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/2
  bfd minimum-interval 150
  bfd multiplier 3
  bfd fast-detect ipv4
  point-to-point
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/3
  point-to-point
  address-family ipv4 unicast
```

Checking the Cisco IOS XR configuration using the **inheritance** keyword in the **show run** command.

```
RP/0/0/CPU0:RT1-XR# show run router isis inheritance
router isis lab
 is-type level-2-only
 net 49.1234.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/0
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/2
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/3
  point-to-point
  address-family ipv4 unicast
   fast-reroute per-prefix level 2
```

If you use details, it indicates which configuration section is inherited from the **group** command.

```
RP/0/0/CPU0:RT1-XR# sh run router isis inheritance detail
router isis lab
 is-type level-2-only
```

```
 net 49.1234.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/0
  point-to-point
  address-family ipv4 unicast
   ## Inherited from group IPFRR
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast∫
   ## Inherited from group IPFRR
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/2
  bfd minimum-interval 150
  bfd multiplier 3
  bfd fast-detect ipv4
  point-to-point
  address-family ipv4 unicast
   ## Inherited from group IPFRR
   fast-reroute per-prefix level 2
  !
 !
 interface GigabitEthernet0/0/0/3
  point-to-point
  address-family ipv4 unicast
   ## Inherited from group IPFRR
   fast-reroute per-prefix level 2
```

Now, let’s check if IP FRR is working in the entire AS 10 network.

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute summary
IS-IS lab IPv4 Unicast FRR summary

                          Critical   High       Medium     Low        Total
                          Priority   Priority   Priority   Priority
Prefixes reachable in L2
  All paths protected     0          0          4          8          12
  Some paths protected    0          0          0          0          0
  Unprotected             0          0          0          0          0
  Protection coverage     0.00%      0.00%      100.00%    100.00%    100.00%
```

```
RT2# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:           Total       Protected     Coverage
   High priority:         0           0             0%
   Normal priority:       14          10            71%
   Total:                 14          10            71%
```

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:           Total       Protected     Coverage
   High priority:         0           0             0%
```

```
   Normal priority:      14         10           71%
   Total:                14         10           71%
```

```
RT4# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        0          0            0%
   Normal priority:      14         10           71%
   Total:                14         10           71%
```

```
RT5# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        0          0            0%
   Normal priority:      14         10           75%
   Total:                14         10           75%
```

It is normal that not all prefixes have protection. It depends on the topology and the algorithm used by IP FRR, for backup path selection.

```
RT2# show isis fast-reroute interfaces
Tag lab - Fast-Reroute Platform Support Information:

   GigabitEthernet5: Protectable: Yes. Usable for repair: Yes
   GigabitEthernet4: Protectable: Yes. Usable for repair: Yes
   GigabitEthernet3: Protectable: Yes. Usable for repair: Yes
   GigabitEthernet2: Protectable: Yes. Usable for repair: Yes
   Loopback0: Protectable: Yes. Usable for repair: Yes
```

```
RT2# show ip route repair-paths
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
i L2     1.1.1.1 [115/10] via 192.168.12.1, 00:27:30, GigabitEthernet2
                       Repair Path: 192.168.25.5, via GigabitEthernet4
                 [RPR][115/20] via 192.168.25.5, 00:27:30, GigabitEthernet4
      2.0.0.0/32 is subnetted, 1 subnets
C        2.2.2.2 is directly connected, Loopback0
      3.0.0.0/32 is subnetted, 1 subnets
i L2     3.3.3.3 [115/20] via 192.168.23.3, 00:27:30, GigabitEthernet5
                       Repair Path: 192.168.12.1, via GigabitEthernet2
                 [RPR][115/30] via 192.168.12.1, 00:27:30, GigabitEthernet2
      4.0.0.0/32 is subnetted, 1 subnets
i L2     4.4.4.4 [115/20] via 192.168.24.4, 00:27:30, GigabitEthernet3
                       Repair Path: 192.168.25.5, via GigabitEthernet4
                 [RPR][115/30] via 192.168.25.5, 00:27:30, GigabitEthernet4
      5.0.0.0/32 is subnetted, 1 subnets
i L2     5.5.5.5 [115/20] via 192.168.25.5, 00:27:30, GigabitEthernet4
                       Repair Path: 192.168.12.1, via GigabitEthernet2
                 [RPR][115/30] via 192.168.12.1, 00:27:30, GigabitEthernet2
      7.0.0.0/32 is subnetted, 1 subnets
B        7.7.7.7 [200/0] via 1.1.1.1, 4d21h
```

 

```
                 [RPR][200/0] via 4.4.4.4, 4d21h
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.12.0/24 is directly connected, GigabitEthernet2
L        192.168.12.2/32 is directly connected, GigabitEthernet2
i L2  192.168.13.0/24 [115/20] via 192.168.23.3, 00:27:30, GigabitEthernet5
                            Repair Path: 192.168.12.1, via GigabitEthernet2

                      [115/20] via 192.168.12.1, 00:27:30, GigabitEthernet2
                            Repair Path: 192.168.23.3, via GigabitEthernet5

i L2  192.168.14.0/24 [115/20] via 192.168.24.4, 00:27:30, GigabitEthernet3
                            Repair Path: 192.168.12.1, via GigabitEthernet2
                      [115/20] via 192.168.12.1, 00:27:30, GigabitEthernet2
                            Repair Path: 192.168.24.4, via GigabitEthernet3

i L2  192.168.15.0/24 [115/20] via 192.168.25.5, 00:27:30, GigabitEthernet4
                            Repair Path: 192.168.12.1, via GigabitEthernet2

                      [115/20] via 192.168.12.1, 00:27:30, GigabitEthernet2
                            Repair Path: 192.168.25.5, via GigabitEthernet4
      192.168.23.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.23.0/24 is directly connected, GigabitEthernet5
L        192.168.23.2/32 is directly connected, GigabitEthernet5
      192.168.24.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.24.0/24 is directly connected, GigabitEthernet3
L        192.168.24.2/32 is directly connected, GigabitEthernet3
      192.168.25.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.25.0/24 is directly connected, GigabitEthernet4
L        192.168.25.2/32 is directly connected, GigabitEthernet4
i L2  192.168.34.0/24 [115/20] via 192.168.24.4, 00:27:30, GigabitEthernet3
                            Repair Path: 192.168.23.3, via GigabitEthernet5

                      [115/20] via 192.168.23.3, 00:27:30, GigabitEthernet5
                            Repair Path: 192.168.24.4, via GigabitEthernet3

i L2  192.168.35.0/24 [115/20] via 192.168.25.5, 00:27:30, GigabitEthernet4
                            Repair Path: 192.168.23.3, via GigabitEthernet5

                      [115/20] via 192.168.23.3, 00:27:30, GigabitEthernet5
                            Repair Path: 192.168.25.5, via GigabitEthernet4

i L2  192.168.45.0/24 [115/20] via 192.168.25.5, 00:27:30, GigabitEthernet4
                            Repair Path: 192.168.24.4, via GigabitEthernet3

                      [115/20] via 192.168.24.4, 00:27:30, GigabitEthernet3
                            Repair Path: 192.168.25.5, via GigabitEthernet4
```

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute

IS-IS lab IPv4 Unicast FRR backups

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       E - EIGRP, A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)
       D - Downstream, LC - Line card disjoint, NP - Node protecting
       P - Primary path, SRLG - SRLG disjoint, TM - Total metric via backup

Maximum parallel path count: 8

L2 2.2.2.2/32 [20/115]
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 30
L2 3.3.3.3/32 [20/115]
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0, Metric: 30
L2 4.4.4.4/32 [20/115]
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 30
```

```
L2 5.5.5.5/32 [20/115]
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 30
C  192.168.12.0/24
     is directly connected, GigabitEthernet0/0/0/0
   L2 RIB backup [35/115]
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 45
C  192.168.13.0/24
     is directly connected, GigabitEthernet0/0/0/1
   L2 RIB backup [20/115]
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0, Metric: 30
C  192.168.14.0/24
     is directly connected, GigabitEthernet0/0/0/2
   L2 RIB backup [20/115]
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 30
C  192.168.15.0/24
     is directly connected, GigabitEthernet0/0/0/3
   L2 RIB backup [20/115]
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 30
L2 192.168.23.0/24 [20/115]
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 20
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 20
L2 192.168.24.0/24 [20/115]
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 20
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 20
L2 192.168.25.0/24 [20/115]
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 20
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0, Metric: 20
L2 192.168.34.0/24 [25/115]
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 25
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 25
L2 192.168.35.0/24 [20/115]
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 20
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0, Metric: 20
L2 192.168.45.0/24 [20/115]
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 20
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.15.5, GigabitEthernet0/0/0/3, RT5, Weight: 0, Metric: 20
```

```
RP/0/0/CPU0:RT1-XR# show cef fast-reroute
Prefix                Next Hop              Interface
-------------------- -------------------- -------------------
2.2.2.2/32          192.168.12.2/32     GigabitEthernet0/0/0/0
                    192.168.14.4/32     GigabitEthernet0/0/0/2 (!)
3.3.3.3/32          192.168.13.3/32     GigabitEthernet0/0/0/1
                    192.168.14.4/32     GigabitEthernet0/0/0/2 (!)
4.4.4.4/32          192.168.12.2/32     GigabitEthernet0/0/0/0 (!)
                    192.168.14.4/32     GigabitEthernet0/0/0/2
5.5.5.5/32          192.168.13.3/32     GigabitEthernet0/0/0/1 (!)
                    192.168.15.5/32     GigabitEthernet0/0/0/3
192.168.23.0/24     192.168.12.2/32     GigabitEthernet0/0/0/0
                    192.168.13.3/32     GigabitEthernet0/0/0/1
192.168.24.0/24     192.168.12.2/32     GigabitEthernet0/0/0/0
                    192.168.14.4/32     GigabitEthernet0/0/0/2
192.168.25.0/24     192.168.12.2/32     GigabitEthernet0/0/0/0
```

```
                     192.168.15.5/32     GigabitEthernet0/0/0/3
192.168.34.0/24      192.168.13.3/32     GigabitEthernet0/0/0/1
                     192.168.14.4/32     GigabitEthernet0/0/0/2
192.168.35.0/24      192.168.13.3/32     GigabitEthernet0/0/0/1
                     192.168.15.5/32     GigabitEthernet0/0/0/3
192.168.45.0/24      192.168.14.4/32     GigabitEthernet0/0/0/2
                     192.168.15.5/32     GigabitEthernet0/0/0/3
```

```
RT2# show ip cef 4.4.4.4/32
4.4.4.4/32
  nexthop 192.168.24.4 GigabitEthernet3 label [implicit-null|24005]()
    repair: attached-nexthop 192.168.12.1 GigabitEthernet2
```

Do the same verification for all other routers and prefixes.

### TROUBLESHOOTING:

We removed the **fast-reroute per-prefix level-2 all** command on RT5. Note that RT2 still has a backup path to reach the Loopback 0 of  RT5.

```
RT2# show ip cef 5.5.5.5/32
5.5.5.5/32
  nexthop 192.168.25.5 GigabitEthernet4 label [18|24006]()
    repair: attached-nexthop 192.168.12.1 GigabitEthernet2
```

The **fast-reroute per-prefix level-2 all** command enables the local router to identify the alternative path for prefixes in the local routing table. Note that, on RT5, none of the prefixes have a repair path.

```
RT5# show ip cef 4.4.4.4/32
4.4.4.4/32
  nexthop 192.168.45.4 GigabitEthernet3
```

```
RT5# show ip cef 2.2.2.2/32
2.2.2.2/32
  nexthop 192.168.25.2 GigabitEthernet5 label [implicit-null|16]()
```

☞ Note that, RT4 still uses RT5 as one of the repair paths to reach 3.3.3.3.

```
RT4# show ip cef 3.3.3.3/32
3.3.3.3/32
  nexthop 192.168.34.3 GigabitEthernet2 [implicit-null|17]()
    repair: attached-nexthop 192.168.45.5 GigabitEthernet6
```

### MORE INFORMATION:

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/xe-3s/iri-xe-3s-book/iri-ip-lfa-frr.html http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/configuration/guide/b_routing_cg42crs/b_routing_cg42crs_chapter_011.html#con_1433839

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/configuration/guide/b_routing_cg42crs/b_routing_cg42crs_chapter_0100.html#con_1393933

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-3/system_management/configuration/guide/b_sysman_cg43crs/b_sysman_cg43crs_chapter_010001.html

 

# 3. Prefix Prioritization

**QUESTION:**

A design engineer plans to deploy MPLS Layer 2 and Layer 3 VPN services on RT1-XR, RT4, and RT5. The BGP next-hop for these services are based on the IPv4 address of the Loopback 0 interface. Configure all routers in AS 10 (RT1-XR, RT2, RT3, RT4, and RT5) to ensure that the IPv4 addresses of the Loopback 0 interfaces of these three routers (RT1-XR, RT4, and RT5) are converged first over the rest of the network prefixes.

**SOLUTION:**

The IS-IS Prefix Prioritization feature defines three groups of prioritization while running the short path algorithm. By default, any /32 prefixes are designated as medium priority and all other prefixes are designated as low priority. To ensure that the Loopback 0 addresses for RT1-XR, RT4, and RT5 will be the first prefixes to be calculated in the shortest path algorithm in all routers of the AS 10 network, those prefixes must have a higher priority as compared to the rest of the prefixes advertised on IS-IS.

**CONFIGURATION:**

Only on RT1-XR, RT4, and RT5 we need to create a prefix tag for Loopback 0.

On RT4 and RT5, the configuration should be:

```
interface Loopback 0
  isis tag 100
!
router isis lab
  ip route priority high tag 100
```

On RT1-XR, the configuration should be:

```
router isis lab
 address-family ipv4 unicast
  spf prefix-priority level 2 high tag 100
 !
 interface Loopback0
  address-family ipv4 unicast
   tag 100
```

On other routers, we just enable the SPF algorithm to run high prioritization on the prefixes tagged as 100.

On RT2 and RT3, the configuration should be:

```
router isis lab
  ip route priority high tag 100
```

Note that, once you configure some prefixes with higher priority, other prefixes that you did not select including /32 prefixes, all become low priority.

### VALIDATION:

Check that, on RT1-XR, RT4, and RT5 you have two high priority prefixes, found as high priority. RT2 and RT3 have three high priority prefixes.

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute summary
IS-IS lab IPv4 Unicast FRR summary

                          Critical   High       Medium     Low        Total
                          Priority   Priority   Priority   Priority
Prefixes reachable in L2
  All paths protected     0          2          0          12         14
  Some paths protected    0          0          0          0          0
  Unprotected             0          0          0          0          0
  Protection coverage     0.00%      100.00%    0.00%      100.00%    100.00%
```

```
RT4# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        2          2            100%
   Normal priority:      12         8            66%
   Total:                14         10           71%
```

```
RT5# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        2          2            100%
   Normal priority:      12         8            66%
   Total:                14         10           71%
```

```
RT2# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        3          3            100%
   Normal priority:      11         7            63%
   Total:                14         10           71%
```

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total      Protected    Coverage
   High priority:        3          3            100%
   Normal priority:      11         7            63%
   Total:                14         10           71%
```

Check the prefixes with high priority on Cisco IOS XR router.

```
RP/0/0/CPU0:RT1-XR# show isis ipv4 unicast fast-reroute detail
IS-IS lab IPv4 Unicast FRR backups

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
       df - level 1 default (closest attached router), su - summary null
       C - connected, S - static, R - RIP, B - BGP, O - OSPF
       E - EIGRP, A - access/subscriber, M - mobile, a - application
       i - IS-IS (redistributed from another instance)
       D - Downstream, LC - Line card disjoint, NP - Node protecting
```

 

```
       P - Primary path, SRLG - SRLG disjoint, TM - Total metric via backup

Maximum parallel path count: 8

L2 2.2.2.2/32 [20/115] low priority
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT2.00-00, 2.2.2.2
L2 3.3.3.3/32 [20/115] low priority
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT3.00-00, 3.3.3.3
L2 4.4.4.4/32 [20/115] high priority
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4 tag 100, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT4.00-00, 4.4.4.4, tag 100
L2 5.5.5.5/32 [20/115] high priority
     via 192.168.15.5, GigabitEthernet0/0/0/3, RT5 tag 100, Weight: 0
       FRR backup via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT5.00-00, 5.5.5.5, tag 100
C  192.168.12.0/24
     is directly connected, GigabitEthernet0/0/0/0
   L2 RIB backup [20/115] low priority
     via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT2.00-00, 2.2.2.2
   L2 adv [10] IS-IS interface
C  192.168.13.0/24
     is directly connected, GigabitEthernet0/0/0/1
   L2 RIB backup [20/115] low priority
     via 192.168.13.3, GigabitEthernet0/0/0/1, RT3, Weight: 0
       FRR backup via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0, Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT3.00-00, 3.3.3.3
   L2 adv [10] IS-IS interface
--snip--
```

```
RT2# show isis rib

IPv4 local RIB for IS-IS process lab

IPV4 unicast topology base (TID 0, TOPOID 0x0) =================
Repair path attributes:
    DS - Downstream, LC - Linecard-Disjoint, NP - Node-Protecting
    PP - Primary-Path, SR - SRLG-Disjoint

1.1.1.1/32
  [115/L2/10] via 192.168.12.1(GigabitEthernet2), from 1.1.1.1, tag 100, LSP[5/5/572]
              (installed)
     repair path: 192.168.25.5(GigabitEthernet4) metric:20 (SR) LSP[5]

3.3.3.3/32
  [115/L2/20] via 192.168.23.3(GigabitEthernet5), from 3.3.3.3, tag 0, LSP[3/3/573]
              (installed)
     repair path: 192.168.12.1(GigabitEthernet2) metric:30 (SR) LSP[3]

4.4.4.4/32
  [115/L2/20] via 192.168.24.4(GigabitEthernet3), from 4.4.4.4, tag 100, LSP[2/2/573]
              (installed)
     repair path: 192.168.25.5(GigabitEthernet4) metric:30 (SR) LSP[2]

5.5.5.5/32
  [115/L2/20] via 192.168.25.5(GigabitEthernet4), from 5.5.5.5, tag 100, LSP[4/4/574]
              (installed)
     repair path: 192.168.12.1(GigabitEthernet2) metric:30 (SR) LSP[4]
```

```
192.168.12.0/24
  [115/L2/20] via 192.168.12.1(GigabitEthernet2), from 1.1.1.1, tag 0, LSP[5/5/572]

192.168.13.0/24
  [115/L2/20] via 192.168.12.1(GigabitEthernet2), from 1.1.1.1, tag 0, LSP[5/5/572]
             (installed)
     repair path: 192.168.23.3(GigabitEthernet5) metric:20 (PP,DS,NP,SR) LSP[3]
  [115/L2/20] via 192.168.23.3(GigabitEthernet5), from 3.3.3.3, tag 0, LSP[3/3/573]
             (installed)
     repair path: 192.168.12.1(GigabitEthernet2) metric:20 (PP,DS,NP,SR) LSP[5]

192.168.14.0/24
  [115/L2/20] via 192.168.24.4(GigabitEthernet3), from 4.4.4.4, tag 0, LSP[2/2/573]
             (installed)
     repair path: 192.168.12.1(GigabitEthernet2) metric:20 (PP,DS,NP,SR) LSP[5]
  [115/L2/20] via 192.168.12.1(GigabitEthernet2), from 1.1.1.1, tag 0, LSP[5/5/572]
             (installed)
     repair path: 192.168.24.4(GigabitEthernet3) metric:20 (PP,DS,NP,SR) LSP[2]
--snip--
```

You can verify the prefixes on the other routers.

### TROUBLESHOOTING:

We removed the **ip route priority high tag 100** command on RT4, to verify the new output and compare to the previous one, before removing this command.

```
RT4# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:          Total        Protected    Coverage
   High priority:        0            0            0%
   Normal priority:      14           10           71%
   Total:                14           10           71%
```

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute summary
IS-IS lab IPv4 Unicast FRR summary

                         Critical   High       Medium     Low        Total
                         Priority   Priority   Priority   Priority
Prefixes reachable in L2
  All paths protected    0          2          0          12         14
  Some paths protected   0          0          0          0          0
  Unprotected            0          0          0          0          0
  Protection coverage    0.00%      100.00%    0.00%      100.00%    100.00%
```

So, RT4 will no longer prioritize Loopback 0 of other routers while running the SPF algorithm.

☞ Note that, other routers still compute the Loopback 0 interface of RT4 as high priority, because, RT4 still tags the Loopback 0 interface.

However, if we remove the **isis tag 100** command under the Loopback 0 interface of RT4, other routers will not prioritize the 4.4.4.4/32 prefix while running the SPF algorithm.

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute summary
IS-IS lab IPv4 Unicast FRR summary

                         Critical   High       Medium     Low        Total
                         Priority   Priority   Priority   Priority
Prefixes reachable in L2
  All paths protected    0          1          0          11         12
```

```
  Some paths protected    0          0          0          0          0
  Unprotected             0          0          0          0          0
  Protection coverage     0.00%      100.00%    0.00%      100.00%    100.00%
```

```
RP/0/0/CPU0:RT1-XR# show isis fast-reroute 4.4.4.4/32 detail
L2 4.4.4.4/32 [20/115] low priority
     via 192.168.14.4, GigabitEthernet0/0/0/2, RT4, Weight: 0
       FRR backup via 192.168.12.2, GigabitEthernet0/0/0/0, RT2, Weight: 0,
Metric: 30
       P: No, TM: 30, LC: No, NP: No, D: No, SRLG: Yes
     src RT4.00-00, 4.4.4.4
```

## MORE INFORMATION:

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/configuration/guide/b_routing_cg42crs/b_routing_cg42crs_chapter_0100.html#con_1396502

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/configuration/guide/b_routing_cg42crs/b_routing_cg42crs_chapter_011.html#task_1263963

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fslocrib.html

 

# 4. Filtering Prefixes that Should Have a Backup Path

**QUESTION:**

Due to memory issues, the design engineer decided to allow pre-computing of backup path only for high priority prefixes on RT3.

**SOLUTION:**

This filter can be specified by using a prefix-list or an access-list.

**CONFIGURATION:**

In this example, we used prefix-list on RT3.

```
ip prefix-list LOOP permit 1.1.1.1/32
ip prefix-list LOOP permit 4.4.4.4/32
ip prefix-list LOOP permit 5.5.5.5/32
!
route-map LFA permit 10
  match ip address prefix-list LOOP
!
router isis lab
  fast-reroute per-prefix level-2 route-map LFA
```

**VALIDATION:**

Output before the change:

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:            Total      Protected    Coverage
   High priority:          3          3            100%
   Normal priority:        11         7            63%
   Total:                  14         10           75%
```

Output after the change:

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:            Total      Protected    Coverage
   High priority:          3          3            100%
   Normal priority:        11         0            0%
   Total:                  14         3            21%
```

 

**TROUBLESHOOTING:**

If the prefix-list is removed by mistake (**no ip prefix-list LOOP** command), the protected backup will be pre-calculated in the same way as if the **route-map** command was not applied in the **fast-reroute per-prefix** command.

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:            Total      Protected    Coverage
   High priority:          3          3            100%
   Normal priority:        11         7            63%
   Total:                  14         10           75%
```

Now, if the **route-map** is removed by mistake (no **route-map LFA**) when there is a route-map reference in the **fast-reroute per-prefix** command**,** then none of the backup paths will be calculated. This is the same as if the route-map is configured and excludes all prefixes for prioritization.

```
RT3# show isis fast-reroute summary
Tag lab:
Microloop Avoidance State: Disabled
IPv4 Fast-Reroute Protection Summary:

 Prefix Counts:            Total      Protected    Coverage
   High priority:          3          0            0%
   Normal priority:        11         0            0%
   Total:                  14         0            0%
```

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/command/irs-cr-book/irs-a1.html#wp5007922600

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-1/routing/command/reference/b_routing_cr41crs/b_routing_cr41crs_chapter_011.html#wp3105567234

 

# 5. Share Risk Link Group (SRLG)

**QUESTION:**

An optical engineer notified the design engineer that some links on the network share a common fiber. This characteristic affects the IP FRR implementation. Enable the share risk link group to avoid the backup path to use the same common fiber that it is used by the primary path.

- On RT1-XR, GigabitEthernet 0/0/0/0 and GigabitEthernet 0/0/0/2 share the same fiber.
- On RT4, GigabitEthernet 4 and GigabitEthernet 6 share the same fiber.

**SOLUTION:**

The Share Risk Link Group (SRLG) feature is used to notify that GigabitEthernet 0/0/0/0 and GigabitEthernet 0/0/0/1 belong to the same group for RT1-XR. Similarly, the SRLG feature should use the same value for GigabitEthernet 2 and GigabitEthernet 3, on RT4. The IS-IS fast-reroute default tie-break should be adjusted to make sure that the SRLG disjoint will be the preferred criterion for backup path selection over the lowest metric value.

**CONFIGURATION:**

We can use any value as long as this value is the same for all interfaces that use the same fiber, and different values for interfaces that use different fiber connections.

The following is the configuration for Cisco IOS XR CLI that is running on RT1-XR.

```
srlg
 interface GigabitEthernet0/0/0/0
  8 value 10
 !
 interface GigabitEthernet0/0/0/1
  8 value 20
 !
 interface GigabitEthernet0/0/0/2
  8 value 10
 !
 interface GigabitEthernet0/0/0/3
  8 value 30
!
router isis lab
 apply-group IPFRR
 is-type level-2-only
 net 49.1234.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
  fast-reroute per-prefix tiebreaker lowest-backup-metric index 50
  fast-reroute per-prefix tiebreaker srlg-disjoint index 40
```

Index can be any value, as long as the **srlg-disjoint index** value is lower than **lowest-backup-metric** value, ensure that the **srlg-disjoint** will be the preferred criterion for backup path selection.

 

The following configuration should be applied on RT4.

```
interface GigabitEthernet 2
  srlg gid 2
!
interface GigabitEthernet 3
  srlg gid 3
!
interface GigabitEthernet 4
  srlg gid 10
!
interface GigabitEthernet 6
  srlg gid 10
!
router isis lab
  fast-reroute tie-break level-1 lowest-backup-path-metric 50
  fast-reroute tie-break level-1 srlg-disjoint 40
```

### VALIDATION:

Before applying the configuration on RT1-XR.

```
RP/0/0/CPU0:RT1-XR# show isis topology
IS-IS lab paths to IPv4 Unicast (Level-2) routers
System Id       Metric  Next-Hop        Interface       SNPA
RT1-XR          --
RT2             10      RT2             Gi0/0/0/0       *PtoP*
RT3             10      RT3             Gi0/0/0/1       *PtoP*
RT4             10      RT4             Gi0/0/0/2       *PtoP*
RT5             10      RT5             Gi0/0/0/3       *PtoP*
```

```
RP/0/0/CPU0:RT1-XR# show route ipv4 4.4.4.4
Routing entry for 4.4.4.4/32
  Known via "isis lab", distance 115, metric 30
  Tag 100, type level-2
  Installed Aug  3 13:32:27.452 for 00:01:44
  Routing Descriptor Blocks
    192.168.12.2, from 4.4.4.4, via GigabitEthernet0/0/0/0, Backup
      Route metric is 30
    192.168.14.4, from 4.4.4.4, via GigabitEthernet0/0/0/2, Protected
      Route metric is 20
  No advertising protos.
```

After applying the configuration on RT1-XR.

```
RP/0/0/CPU0:RT1-XR# show route ipv4 4.4.4.4
Routing entry for 4.4.4.4/32
  Known via "isis lab", distance 115, metric 30
  Tag 100, type level-2
  Installed Aug  3 13:35:14.370 for 00:00:48
  Routing Descriptor Blocks
    192.168.13.3, from 4.4.4.4, via GigabitEthernet0/0/0/1, Backup
      Route metric is 30
    192.168.14.4, from 4.4.4.4, via GigabitEthernet0/0/0/2, Protected
      Route metric is 20
  No advertising protos.
```

 

Before applying the configuration on RT4.

```
RT4# show isis topology
Tag lab:

IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
RT1-XR               10         RT1-XR               Gi4         fa16.3e6d.4baf
RT2                  10         RT2                  Gi3         fa16.3ef8.80d1
RT3                  10         RT3                  Gi2         fa16.3e22.3a17
RT4                  --
RT5                  10         RT5                  Gi6         fa16.3e04.3501
```

```
RT4# show ip route 1.1.1.1
Routing entry for 1.1.1.1/32
  Known via "isis", distance 115, metric 20
  Tag 100, type level-2
  Redistributing via isis lab
  Last update from 192.168.24.2 on GigabitEthernet3, 00:01:15 ago
  Routing Descriptor Blocks:
  * 192.168.14.1, from 1.1.1.1, 01:26:32 ago, via GigabitEthernet4
      Route metric is 10, traffic share count is 1
      Route tag 100
      Repair Path: 192.168.45.5, via GigabitEthernet6
```

After applying the configuration on RT4.

```
RT4# show ip route 1.1.1.1
Routing entry for 1.1.1.1/32
  Known via "isis", distance 115, metric 20
  Tag 100, type level-2
  Redistributing via isis lab
  Last update from 192.168.24.2 on GigabitEthernet3, 00:00:24 ago
  Routing Descriptor Blocks:
* 192.168.14.1, from 1.1.1.1, 00:07:38 ago, via GigabitEthernet4
      Route metric is 10, traffic share count is 1
      Route tag 100
      Repair Path: 192.168.24.2, via GigabitEthernet3
```

**TROUBLESHOOTING:**

Now, practice the lab changing and removing some parameters in the configuration and see how it impacts the result (one at the time). Also, remove the **fast-reroute tiebreak** command and see what happens.

Lastly, change the metric of the backup paths and identify that the higher metric is chosen because the lowest metric path uses the same fiber as the primary path.

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/routers/ncs6000/software/segment-routing/configuration/guide/b-segment-routing-cg-ncs6k/b-segment-routing-cg-ncs6k_chapter_01000.pdf

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/mpls/command/reference/b_mpls_cr42crs/b_mpls_cr42crs_chapter_011.html#wp2881578553

http://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/mpls/configuration/guide/b_mpls_cg42crs/b_mpls_cg42crs_chapter_0100.html#task_CEBB3334EFF642D6B6FD604CCC6F539E

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-3s/iro-lfa-frr-xe.html

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/15-s/irs-15-s-book/irs-rmte-lfa-frr.html

 

# 6. MPLS LDP Session Protection

**QUESTION:**

Customize the LDP configuration to ensure that when a link between the two LSRs fails, the LDP session stays up, consequently the LDP bindings for prefixes do not need to be relearned. Implement this optimization for the LDP session between RT1-XR and RT5.

**SOLUTION:**

MPLS LDP session protection uses LDP Targeted Hellos to protect LDP sessions. As long as the LDP peer is reachable through IP, when the direct link between the two peers fails, the LDP session still stays up.

**CONFIGURATION:**

The following configuration is to be applied on RT1-XR.

```
ipv4 access-list RT5
 10 permit ipv4 host 5.5.5.5 any

mpls ldp
 router-id 1.1.1.1
 session protection for RT5
```

The following configuration is to be applied on RT5.

```
access-list 6 permit host 1.1.1.1
!
mpls ldp session protection for 6
```

**VALIDATION:**

Cisco IOS-XR has a specific show command to verify the MPLS LDP session protection.

```
RP/0/0/CPU0:RT1-XR# show mpls ldp neighbor sp

Peer LDP Identifier: 5.5.5.5:0
  TCP connection: 5.5.5.5:45628 - 1.1.1.1:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 3293/3305; Downstream-Unsolicited
  Up time: 1d23h
  LDP Discovery Sources:
    IPv4: (2)
      GigabitEthernet0/0/0/3
      Targeted Hello (1.1.1.1 -> 5.5.5.5, active)
    IPv6: (0)
  Addresses bound to this peer:
    IPv4: (4)
      5.5.5.5        192.168.15.5   192.168.25.5   192.168.35.5
      192.168.45.5
    IPv6: (0)
```

 

On RT5 router, you must use the **detail** keyword to verify the session protection information.

```
RT5# show mpls ldp neighbor 1.1.1.1 detail
Peer LDP Ident: 1.1.1.1:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 1.1.1.1.646 - 5.5.5.5.45628
        Password: not required, none, in use
        State: Oper; Msgs sent/rcvd: 3318/3307; Downstream; Last TIB rev sent 26
        Up time: 2d00h; UID: 3; Peer Id 2
        LDP discovery sources:
          GigabitEthernet2; Src IP addr: 192.168.15.1
            holdtime: 15000 ms, hello interval: 5000 ms
          Targeted Hello 5.5.5.5 -> 1.1.1.1, active, passive;
            holdtime: infinite, hello interval: 10000 ms
        Addresses bound to peer LDP Ident:
          192.168.12.1    192.168.13.1    192.168.14.1    192.168.15.1
          172.16.1.1      1.1.1.1
        Peer holdtime: 180000 ms; KA interval: 60000 ms; Peer state: estab
        Clients: Dir Adj Client
        LDP Session Protection enabled, state: Ready
            acl: 6, duration: 86400 seconds
        NSR: Not Ready
        Capabilities Sent:
          [ICCP (type 0x0405) MajVer 1 MinVer 0]
          [Dynamic Announcement (0x0506)]
          [mLDP Point-to-Multipoint (0x0508)]
          [mLDP Multipoint-to-Multipoint (0x0509)]
          [Typed Wildcard (0x050B)]
        Capabilities Received:
          [mLDP Point-to-Multipoint (0x0508)]
          [mLDP Multipoint-to-Multipoint (0x0509)]
          [Typed Wildcard (0x050B)]
```

If we shut down the link between RT5 and RT1-XR, you will notice that the MPLS LDP session stays up.

```
RT5(config)# interface gigabitethernet 2
RT5(config-if)# shut
*Aug  3 15:48:39.219: %LDP-5-SP: 1.1.1.1:0: session hold up initiated
RT5(config-if)# end
*Aug  3 15:48:41.210: %LINK-5-CHANGED: Interface GigabitEthernet2, changed state
to administratively down
*Aug  3 15:48:41.520: %SYS-5-CONFIG I: Configured from console by console
*Aug  3 15:48:42.210: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet2, changed state to down
*Aug  3 15:48:44.773: %BFD-6-BFD SESS DESTROYED: BFD-SYSLOG:
bfd session destroyed,  handle:1 neigh proc:FRR, handle:1 act
RT5# sh mpls ldp neighbor 1.1.1.1
     Peer LDP Ident: 1.1.1.1:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 1.1.1.1.646 - 5.5.5.5.45628
        State: Oper; Msgs sent/rcvd: 3323/3310; Downstream
        Up time: 2d00h
        LDP discovery sources:
          Targeted Hello 5.5.5.5 -> 1.1.1.1, active, passive
        Addresses bound to peer LDP Ident:
--snip--
```

### TROUBLESHOOTING:

Remove the configuration on RT5 and verify the MPLS LDP session protection on RT1-XR side with the **show mpls ldp neighbor sp** command. It might give the impression that the LDP session protection is activated.

```
RP/0/0/CPU0:RT1-XR# show mpls ldp neighbor sp
Peer LDP Identifier: 5.5.5.5:0
  TCP connection: 5.5.5.5:22975 - 1.1.1.1:646
  Graceful Restart: No
```

```
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 23/23; Downstream-Unsolicited
  Up time: 00:06:10
  LDP Discovery Sources:
    IPv4: (2)
      GigabitEthernet0/0/0/3
      Targeted Hello (1.1.1.1 -> 5.5.5.5, active)
    IPv6: (0)
  Addresses bound to this peer:
--snip--
```

However, if we shut down the interface between RT5 and RT1-XR, this LDP session drops.

```
RT5(config)# interface GigabitEthernet 2
RT5(config-if)# shut
*Aug  3 16:21:41.068: %LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 (3) is DOWN (Session
Protection disabled targeted session)
*Aug  3 16:21:43.056: %LINK-5-CHANGED: Interface GigabitEthernet2, changed state
to administratively down
*Aug  3 16:21:44.057: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet2, changed state to down
```

```
RP/0/0/CPU0:RT1-XR# show mpls ldp neighbor sp

RP/0/0/CPU0:RT1-XR#
```

In the Cisco IOS output, you need to look for the sentence "`LDP Session Protection enabled, state: Ready`" to ensure that LDP session protection is activated and functional.

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-3s/mp-ldp-xe-3s-book/mp-ldp-sessn-prot.pdf

http://www.cisco.com/c/en/us/td/docs/ios_xr_sw/iosxr_r3-7/mpls/configuration/guide/gc37ldp.html#wp1148788

# 7. LDP IGP Sync

**QUESTION:**

An operations engineer identified packet loss when the IS-IS adjacency is just established, the router begins forwarding packets using the new adjacency before the LDP label exchange completes between the peers on that link. Optimize the AS 10 network (RT1-XR, RT2, RT3, RT4, and RT5) that provides a means to synchronize LDP and IS-IS to minimize the MPLS packet loss.

**SOLUTION:**

The MPLS LDP-IGP synchronization feature:

- Provides a means to synchronize LDP and IGPs to minimize MPLS packet loss.
- Prevents MPLS packet loss due to synchronization conflicts.

**CONFIGURATION:**

The following is the configuration for Cisco IOS devices on AS 10 (RT2, RT3, RT4, and RT5).

```
router isis lab
  mpls ldp sync
```

For Cisco IOS XR we have an option to use group. This configuration can be applied on RT1-XR. The following solution uses the **group** command. This group command enables LDP IGP Sync for all IS-IS enabled interfaces using a regular expression.

```
group LDP-IGP-SYNC
 router isis '.*'
  interface 'Giga.*'
   address-family ipv4 unicast
    mpls ldp sync level 2
!
apply-group LDP-IGP-SYNC
```

**VALIDATION:**

```
RP/0/0/CPU0:RT1-XR# show mpls ldp igp sync
GigabitEthernet0/0/0/0:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
    Peers:
      2.2.2.2:0

GigabitEthernet0/0/0/1:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
    Peers:
      3.3.3.3:0

GigabitEthernet0/0/0/2:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
    Peers:
      4.4.4.4:0
```

```
GigabitEthernet0/0/0/3:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
    Peers:
      5.5.5.5:0
```

```
RT4# show isis mpls ldp
Interface: GigabitEthernet6; ISIS tag lab enabled
 ISIS is UP on interface
 AUTOCONFIG Information :
   LDP enabled: NO
 SYNC Information :
   Required: YES
   Achieved: YES
   IGP Delay: NO
   Holddown time: Infinite
   State: SYNC achieved

Interface: GigabitEthernet4; ISIS tag lab enabled
 ISIS is UP on interface
 AUTOCONFIG Information :
   LDP enabled: NO
 SYNC Information :
   Required: YES
   Achieved: YES
   IGP Delay: NO
   Holddown time: Infinite
   State: SYNC achieved

Interface: GigabitEthernet3; ISIS tag lab enabled
 ISIS is UP on interface
 AUTOCONFIG Information :
   LDP enabled: NO
 SYNC Information :
   Required: YES
   Achieved: YES
   IGP Delay: NO
   Holddown time: Infinite
   State: SYNC achieved

Interface: GigabitEthernet2; ISIS tag lab enabled
 ISIS is UP on interface
 AUTOCONFIG Information :
   LDP enabled: NO
 SYNC Information :
   Required: YES
   Achieved: YES
   IGP Delay: NO
   Holddown time: Infinite
   State: SYNC achieved

Interface: Loopback0; ISIS tag lab enabled
 ISIS is UP on interface
 AUTOCONFIG Information :
   LDP enabled: NO
 SYNC Information :
   Required: NO
```

## MORE INFORMATION:

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fsldpsyn.html
http://www.cisco.com/c/en/us/td/docs/ios_xr_sw/iosxr_r3-7/mpls/configuration/guide/gc37ldp.html#wp1166228

# 8. BGP Advertise External Path

**QUESTION:**

RT1-XR is set as preferred exit point for AS 10. However, to improve BGP fast convergence, route reflector RT2 and RT3 should receive the information from RT4 as well. Configure AS 10 so that it meets this requirement.

**Note**: You cannot change any configuration on RT2 and RT3.

**SOLUTION:**

Enabling BGP advertise External Path on RT4 forces this router to advertise the external prefixes to the route-reflector even though the best path is an internal path (RT1-XR).

**CONFIGURATION:**

We just need to configure RT4. This is the configuration to be applied on RT4.

```
! RT4
router bgp 10
address-family ipv4
  bgp advertise-best-external
```

**VALIDATION:**

Before implementing the configuration on RT4.

```
RT2# show bgp ipv4 unicast 7.7.7.7/32
BGP routing table entry for 7.7.7.7/32, version 6
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     1          2
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    1.1.1.1 (metric 10) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
```

```
RT3# show bgp ipv4 unicast 7.7.7.7/32
BGP routing table entry for 7.7.7.7/32, version 6
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1          2
  Refresh Epoch 1
  65000, (Received from a RR-client)
    1.1.1.1 (metric 10) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal
```

 

```
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0
```

```
RT4# show bgp ipv4 unicast 7.7.7.7
Paths: (3 available, best #2, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  65000
    172.16.4.7 from 172.16.4.7 (7.7.7.7)
      Origin IGP, metric 0, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
```

```
RT4# show bgp ipv4 unicast neighbor 2.2.2.2 advertised-routes

Total number of prefixes 0
```

After implementing the configuration on RT4.

☞ Note that neighbor 4.4.4.4 is advertising the prefix to the route reflectors.

```
RT2# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 4
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
     1          6
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    4.4.4.4 (metric 20) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    1.1.1.1 (metric 10) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
```

```
RT3# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 4
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
     1          9
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    4.4.4.4 (metric 20) from 4.4.4.4 (4.4.4.4)
```

```
      Origin IGP, metric 0, localpref 100, valid, internal
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    1.1.1.1 (metric 10) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
```

```
RT4# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 6
Paths: (3 available, best #2, table default)
  Advertise-best-external
  Advertised to update-groups:
     1          2
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
   rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  65000
    172.16.4.7 from 172.16.4.7 (7.7.7.7)
      Origin IGP, metric 0, localpref 100, valid, external, backup/repair,
advertise-best-external , recursive-via-connected
      rx pathid: 0, tx pathid: 0
```

However, router reflectors still advertise only the best path to RT5.

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-best-external.html

http://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-1/routing/configuration/guide/routing_cg41xr12k_chapter1.html#task_1724299

# 9. BGP Diverse Path using a Shadow Route Reflector

**QUESTION:**

Configure AS 10 to ensure that IPv4 BGP fast convergence is achieved for both link and node failures by leveraging diverse paths.

**Note**: Configuration of the BGP additional-path feature is not allowed on any route-reflector. The multipath feature is not allowed either.

**SOLUTION:**

Path Diversity can be achieved by:

- BGP Diverse Path: uses a shadow route reflector
- BGP Add-Path: uses Path-ID, unique per prefix

To meet the requirement of this scenario we need to use one of the route reflectors as a shadow router. Shadow route reflector means to advertise a second-best path and *not* the best path information.

**CONFIGURATION:**

We will use RT3 as the shadow route reflector, and this is the configuration to be applied on RT3.

```
router bgp 10
 address-family ipv4
  bgp additional-paths select backup
  neighbor 1.1.1.1 advertise diverse-path backup
  neighbor 4.4.4.4 advertise diverse-path backup
  neighbor 5.5.5.5 advertise diverse-path backup
 exit-address-family
```

**VALIDATION:**

Check that, RT3 is no longer advertising the best-path, instead it is advertising the 2nd best.

```
RT3# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 5
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
     9          10
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    4.4.4.4 (metric 20) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65000, (Received from a RR-client)
    1.1.1.1 (metric 10) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
```

```
      rx pathid: 0, tx pathid: 0x0
```

```
RT3# show bgp ipv4 unicast neighbor 5.5.5.5 advertised-routes
BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *bia7.7.7.7/32       4.4.4.4                  0    100      0 65000 i

Total number of prefixes 1
```

☞ Note that, RT5 receives the information regards to the 7.7.7.7 prefix with BGP next-hop as 4.4.4.4. However, the path is not marked as 'backup/repair'.

Observe the new BGP attributes information: **rx pathid** and **tx pathid**.

```
RT5# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 5
Paths: (2 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 25) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 4.4.4.4, Cluster list: 3.3.3.3
      rx pathid: 0, tx pathid: 0
```

If we would like RT5 to install the backup path to increase the fast convergence in the RIB and FIB, we need to use the BGP additional-paths install feature. Find more information about the BGP additional-paths feature in the next section.

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg_diverse_path.pdf

http://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-1/routing/configuration/guide/routing_cg41xr12k_chapter1.html#con_1723711

# 10. BGP Additional-Path

**QUESTION:**

Configure the routers in AS 10 to ensure that IPv4 BGP fast convergence is achieved for both link and node failures by leveraging the redundant paths at the Core and Edge of the network.

**Note**: Configuration of BGP diverse-path and multipath features are *not* allowed.

**SOLUTION:**

We need to remove the solution of task 8 and 9 to understand better how BGP additional-path works and also to understand the differences between the BGP diverse-path and the BGP additional-path feature. The BGP Additional-Path feature enables the BGP speakers to advertise more than one path for the same prefix. This feature also gives an option if you would like to add an alternative path in the RIB to improve the fast convergence by using the **install** parameter in the **additional-paths** command.

**CONFIGURATION:**

The following output is the configuration for the Cisco IOS-XR route-reflector-client, RT1-XR.

```
route-policy BGP-ADD-PATH
  set path-selection backup 1 advertise install
end-policy
!
router bgp 10
 address-family ipv4 unicast
  additional-paths receive
  additional-paths send
  additional-paths selection route-policy BGP-ADD-PATH
```

The following output is the configuration for the Cisco IOS route-reflector-clients: RT4 and RT5.

```
router bgp 10
 address-family ipv4
  bgp additional-paths select all
  bgp additional-paths send receive
  bgp additional-paths install
  neighbor 2.2.2.2 advertise additional-paths all
  neighbor 3.3.3.3 advertise additional-paths all
```

The following output is the configuration for the route-reflector. The following should be applied on RT2.

```
router bgp 10
 address-family ipv4
  bgp additional-paths select all
  bgp additional-paths send receive
  bgp additional-paths install
  neighbor 1.1.1.1 advertise additional-paths all
  neighbor 3.3.3.3 advertise additional-paths all
  neighbor 4.4.4.4 advertise additional-paths all
  neighbor 5.5.5.5 advertise additional-paths all
 exit-address-family
```

The following output should be applied on RT3.

```
router bgp 10
 address-family ipv4
  bgp additional-paths select all
  bgp additional-paths send receive
  bgp additional-paths install
  neighbor 1.1.1.1 advertise additional-paths all
  neighbor 2.2.2.2 advertise additional-paths all
  neighbor 4.4.4.4 advertise additional-paths all
  neighbor 5.5.5.5 advertise additional-paths all
 exit-address-family
```

**VALIDATION:**

☞ Note that, now both router reflectors advertise two paths for 7.7.7.7/32 prefix: One as BGP next-hop 1.1.1.1 and another as BGP next-hop 4.4.4.4

```
RT5# show bgp ipv4 unicast
BGP table version is 10, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
 * ia                 4.4.4.4                  0    100      0 65000 i
 *>i                  1.1.1.1                  0    200      0 65000 i
 *bi                  4.4.4.4                  0    100      0 65000 i
```

See if you understand the information provided by the two new BGP attributes: **rx pathid** and **tx pathid**.

```
RT5# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 11
Paths: (4 available, best #3, table default)
  Additional-path-install
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0x0, tx pathid: 0
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, all
      Originator: 4.4.4.4, Cluster list: 3.3.3.3, 2.2.2.2
      rx pathid: 0x2, tx pathid: 0x2
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0x0, tx pathid: 0x0
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
```

```
    rx pathid: 0x1, tx pathid: 0
```

☞ Note that, the backup path is in the CEF table as well.

```
RT5# show ip cef 7.7.7.7/32 internal
7.7.7.7/32, epoch 2, flags [rnolbl, rlbls], RIB[B], refcnt 6, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00018000
    Broker: linked, distributed at 4th priority
  ifnums:
    GigabitEthernet2(8): 192.168.15.1
    GigabitEthernet3(9): 192.168.45.4
    GigabitEthernet4(10): 192.168.35.3
    GigabitEthernet5(11): 192.168.25.2
  path list 7FB212570B88, 3 locks, per-destination, flags 0x269 [shble, rif, rcrsv, hwcn, bgp]
    path 7FB21890F6D0, share 1/1, type recursive, for IPv4
      recursive via 1.1.1.1[IPv4:Default], fib 7FB1FBE644F0, 1 terminal fib, v4:Default:1.1.1.1/32
      path list 7FB212570AE8, 3 locks, per-destination, flags 0x49 [shble, rif, hwcn]
          path 7FB21890F580, share 1/1, type attached nexthop, for IPv4, flags [has-rpr]
            MPLS short path extensions: [none] MOI flags = 0x20 label implicit-null
            nexthop 192.168.15.1 GigabitEthernet2 label [implicit-null|25](), IP adj out of
GigabitEthernet2, addr 192.168.15.1 7FB20D5F3458
              repair: attached-nexthop 192.168.25.2 GigabitEthernet5 (7FB21890F628)
          path 7FB21890F628, share 1/1, type attached nexthop, for IPv4, flags [rpr, rpr-only]
            nexthop 192.168.25.2 GigabitEthernet5, repair, IP adj out of GigabitEthernet5, addr
192.168.25.2 7FB20D5F3828
    path 7FB21890F820, share 1/1, type recursive, for IPv4, flags [rpr]
      recursive via 4.4.4.4[IPv4:Default], repair, fib 7FB1FBE64DF0, 1 terminal fib,
v4:Default:4.4.4.4/32
      path list 7FB212571128, 3 locks, per-destination, flags 0x49 [shble, rif, hwcn]
          path 7FB21890F388, share 1/1, type attached nexthop, for IPv4, flags [has-rpr]
            MPLS short path extensions: [none] MOI flags = 0x20 label implicit-null
            nexthop 192.168.45.4 GigabitEthernet3 label [implicit-null|17](), IP adj out of
GigabitEthernet3, addr 192.168.45.4 7FB20D5F3FC8
              repair: attached-nexthop 192.168.35.3 GigabitEthernet4 (7FB218910150)
          path 7FB218910150, share 1/1, type attached nexthop, for IPv4, flags [rpr, rpr-only]
            nexthop 192.168.35.3 GigabitEthernet4, repair, IP adj out of GigabitEthernet4, addr
192.168.35.3 7FB20D5F3BF8
  output chain:
    label [implicit-null|25]()
    FRR Primary (0x7FB1FBFC54E8)
      <primary: IP adj out of GigabitEthernet2, addr 192.168.15.1 7FB20D5F3458>
      <repair:  TAG adj out of GigabitEthernet5, addr 192.168.25.2 7FB20D5F3640>
```

**TROUBLESHOOTING:**

On RT5 we removed the command **bgp additional-path install.** Now, verify the BGP database and compare with the previous output when everything was configured correctly.

```
RT5# show bgp ipv4 unicast
BGP table version is 12, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
 * ia                 4.4.4.4                  0    100      0 65000 i
 *>i                  1.1.1.1                  0    200      0 65000 i
 * i                  4.4.4.4                  0    100      0 65000 i
```

☞ Note that, you can't find the flag **b** in the output above and the **backup** word information in the **detail** output command displayed below.

```
RT5# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 12
Paths: (4 available, best #3, table default)
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0x0, tx pathid: 0
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, all
      Originator: 4.4.4.4, Cluster list: 3.3.3.3, 2.2.2.2
      rx pathid: 0x2, tx pathid: 0x2
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0x0, tx pathid: 0x0
  Path not advertised to any peer
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx pathid: 0x1, tx pathid: 0
```

☞ Note that, in the CEF output you cannot find the backup path with 4.4.4.4 as the BGP next-hop.

```
RT5# show ip cef 7.7.7.7/32 internal
7.7.7.7/32, epoch 2, flags [rnolbl, rlbls], RIB[B], refcnt 6, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00018000
    Broker: linked, distributed at 4th priority
  ifnums:
    GigabitEthernet2(8): 192.168.15.1
    GigabitEthernet5(11): 192.168.25.2
  path list 7FB212570FE8, 3 locks, per-destination, flags 0x269 [shble, rif, rcrsv, hwcn, bgp]
    path 7FB21890FF58, share 1/1, type recursive, for IPv4
      recursive via 1.1.1.1[IPv4:Default], fib 7FB1FBE644F0, 1 terminal fib, v4:Default:1.1.1.1/32
      path list 7FB212570AE8, 3 locks, per-destination, flags 0x49 [shble, rif, hwcn]
          path 7FB21890F580, share 1/1, type attached nexthop, for IPv4, flags [has-rpr]
            MPLS short path extensions: [none] MOI flags = 0x20 label implicit-null
            nexthop 192.168.15.1 GigabitEthernet2 label [implicit-null|25](), IP adj out of
GigabitEthernet2, addr 192.168.15.1 7FB20D5F3458
              repair: attached-nexthop 192.168.25.2 GigabitEthernet5 (7FB21890F628)
          path 7FB21890F628, share 1/1, type attached nexthop, for IPv4, flags [rpr, rpr-only]
            nexthop 192.168.25.2 GigabitEthernet5, repair, IP adj out of GigabitEthernet5, addr
192.168.25.2 7FB20D5F3828
  output chain:
    label [implicit-null|25]()
    FRR Primary (0x7FB1FBFC54E8)
      <primary: IP adj out of GigabitEthernet2, addr 192.168.15.1 7FB20D5F3458>
      <repair:  TAG adj out of GigabitEthernet5, addr 192.168.25.2 7FB20D5F3640>
```

Now, we removed also the **bgp additional-path send receive** command on RT5.

RT5 no longer receives (accepts to be more precise) the information with regards to BGP next-hop 4.4.4.4.

```
RT5# show bgp ipv4 unicast
BGP table version is 16, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
 * i                  1.1.1.1                  0    200      0 65000 i
```

```
RT5# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 16
Paths: (2 available, best #1, table default)
  Path not advertised to any peer
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0, tx pathid: 0x0
  Path not advertised to any peer
  Refresh Epoch 1
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0, tx pathid: 0
```

☞ Note that, this is negotiated between peers, so if you check on RT2, you will notice that RT2 sends the backup path to other peers but not to 5.5.5.5. This is because in the BGP negotiation, RT5 did not send the message for the BGP Additional-path.

```
RT2# sh bgp ipv4 unicast neighbor 5.5.5.5 advertised-routes
BGP table version is 15, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
```

```
RT2# show bgp ipv4 unicast neighbor 4.4.4.4 advertised-routes
BGP table version is 15, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
 *bia7.7.7.7/32       4.4.4.4                  0    100      0 65000 i

Total number of prefixes 2
```

Now, we returned all the commands we removed and then we removed only the **bgp additional-paths select all** command on RT5. Just to clarify what is the  final configuration on RT5 at this point, it follows the BGP configuration for RT5.

```
RT5# sh run | section address-family ipv4
 address-family ipv4
  bgp additional-paths send receive
  bgp additional-paths install
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.2 inherit peer-policy BGP
  neighbor 2.2.2.2 advertise additional-paths all
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 inherit peer-policy BGP
  neighbor 3.3.3.3 advertise additional-paths all
```

☞ Note that, now we do not have the flag **a,** neither the **all** word in the **show bgp ipv4 unicast 7.7.7.7** output command.

```
RT5# show bgp ipv4 unicast
BGP table version is 21, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 7.7.7.7/32       1.1.1.1                  0    200      0 65000 i
 * i                  4.4.4.4                  0    100      0 65000 i
 *>i                  1.1.1.1                  0    200      0 65000 i
 *bi                  4.4.4.4                  0    100      0 65000 i
```

```
RT5# show bgp ipv4 unicast 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 21
Paths: (4 available, best #3, table default)
  Additional-path-install
  Not advertised to any peer
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 200, valid, internal
      Originator: 1.1.1.1, Cluster list: 3.3.3.3
      rx pathid: 0x0, tx pathid: 0
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 4.4.4.4, Cluster list: 3.3.3.3, 2.2.2.2
      rx pathid: 0x2, tx pathid: 0
  Refresh Epoch 2
  65000
    1.1.1.1 (metric 10) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 1.1.1.1, Cluster list: 2.2.2.2
      rx pathid: 0x0, tx pathid: 0x0
  Refresh Epoch 2
  65000
    4.4.4.4 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx pathid: 0x1, tx pathid: 0
```

 

**MORE INFORMATION:**

http://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-1/routing/configuration/guide/routing_cg41xr12k_chapter1.html#task_6CA10985EC294186A710FAD19A062605

https://www.cisco.com/c/en/us/td/docs/routers/crs/software/crs_r4-2/routing/configuration/guide/b_routing_cg42crs/b_routing_cg42crs_chapter_01.html#concept_77EE033C2F0C4BDDB8423C25FA71E3F9

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-additional-paths.html

 