

LFA Policies Using OSPF as IGP

In This Chapter

This section provides information about LFA policies using OSPF as IGP.

Topics in this section include:

- [Applicability on page 1786](#)
- [Overview on page 1787](#)
- [Configuration on page 1789](#)
- [Conclusion on page 1802](#)

Applicability

Loop Free Alternate (LFA) policies is a local control plane feature. The functionality is limited to the hardware supported by LDP Fast ReRoute (FRR) and IP FRR:

- LDP FRR is supported on the 7750 SR-7/12, 7450 ESS-6/6v/7/12 in all chassis modes, on the 7450 ESS-6/6v/7/12 in mixed-mode and on the 7950 XRS-16/20/40. It is also supported on the 7750 SR-c4/12 and 7710 SR-c4/c12 platforms.
- IP FRR is supported on the 7750 SR-7/12 in chassis mode D, on the 7450 ESS-6/6v/7/12 in chassis mode D with or without mixed-mode. It is also supported on the 7750 SR-c4/12 platforms.

This configuration was tested on release 12.0.R4.

Overview

When multiple LFAs exist, RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*, chooses the selection of the LFA providing the best coverage of the failure cases. In general, this means that node LFA has preference above link LFA. In some deployments, however, this can lead to suboptimal LFA. For example an aggregation router (typically using lower bandwidth links) protecting a core node/link (typically using high bandwidth links) is potentially undesirable.

For this reason, the operator wants to have more control in the LFA next-hop selection algorithm. This is achieved by the introduction of LFA Shortest Path First (SPF) policies.

LFA policies can work in combination with IP FRR and/or LDP FRR.

Implementation

The 7x50 LFA policy implementation is built around the concept of route next-hop (NH) templates which are applied to IP interfaces. A route-next-hop template specifies criteria which influence the selection of an LFA backup NH for either:

- a set of prefixes in a prefix-list or
- a set of prefixes which resolve to a specific primary NH

Refer to <http://tools.ietf.org/html/draft-litkowski-rtgwg-lfa-manageability> for further information. Two powerful methods which can be used as criteria inside a route-next-hop template are IP admin-groups and IP Shared Risk Link Group (SRLG). IP SRLG and IP admin-group criteria are applied before running the LFA NH algorithm. IP Admin-groups and IP SRLG work in a similar way as the well-known MPLS admin-groups and MPLS SRLG.

For example, when one or more IP admin-groups/SRLG are applied to an IP interface, the same MPLS admin-group/SRLG rules apply:

- IP interfaces which do not include one or more of the admin-groups in the **include** statements are pruned before computing the LFA next-hop.
- IP interfaces which belong to admin-groups which have been explicitly excluded using the **exclude** statement are pruned before computing the LFA next-hop.
- IP interfaces which belong to the SRLGs used by the primary NH of a prefix are pruned before computing the LFA next-hop.

For compatibility reasons with the existing MPLS, admin-groups and SRLG, a single set of admin-groups and SRLGs are defined within the **configure router if-attribute** context from 12.0.R1 onward. Configuration of admin-groups and SRLGs in the **configure router mpls** context is deprecated from this release onwards.

Implementation

Once one or more admin-groups/SRLGs have been defined, it is possible to apply them on an MPLS interface and/or an IP interface.

In the current implementation IP admin-groups/SRLGs are locally significant, meaning they are not advertised by the IGP.

Keep in mind that the well-known MPLS admin-groups/SRLGs are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in the IGP protocol.

Other selection criteria which can be configured inside a route-next-hop template are protection type preference and NH type preference. More details on these parameters are provided later in this example.

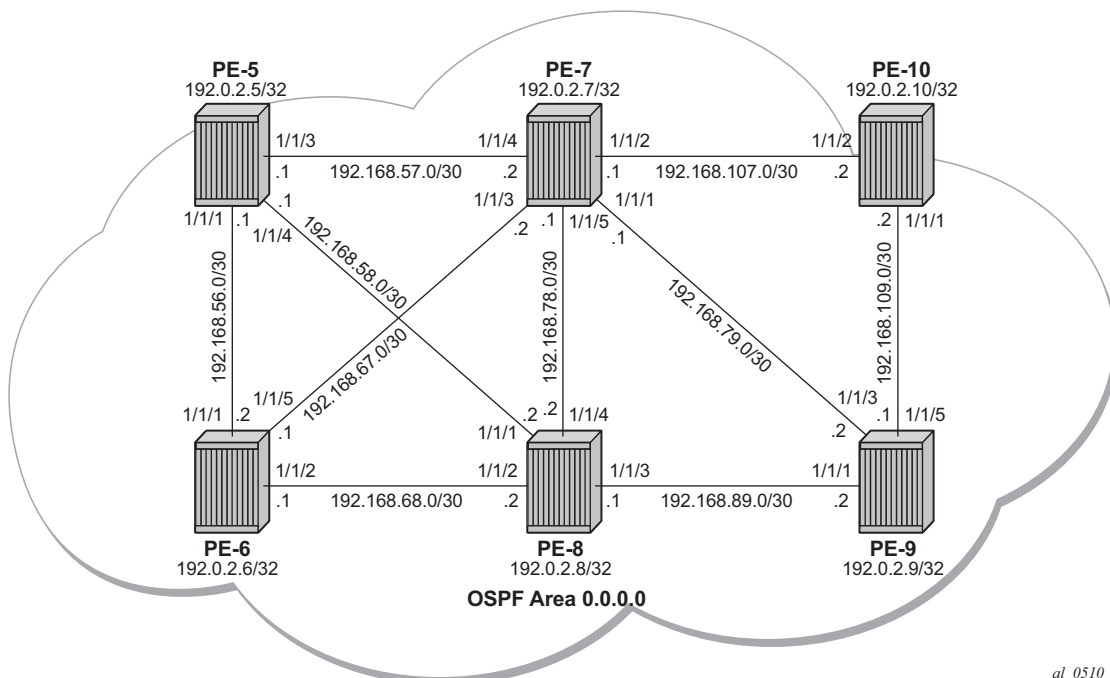


Figure 279: Network Topology

al_0510

Configuration

Step 1. Configuring an IP/MPLS network with LDP FRR enabled on PE-7.

Since the focus is not on how to setup an IP/MPLS network, only summary bullets are provided.

- The system and IP interface addresses are configured according to [Figure 279](#).
- OSPF area 0 is selected as the interior gateway protocol (IGP) to distribute routing information between all PEs. All OSPF interfaces are setup as **type point-to-point** to avoid running the DR/BDR election process.
- Enable link LDP on all interfaces. This establishes a full mesh of LDP LSPs between all PEs' system interfaces. As an example, the tunnel-table on PE-7 looks like this:

```
*A:PE-7# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId  Pref   Nexthop      Metric
-----
192.0.2.5/32         ldp   MPLS    -        9      192.168.57.1  1000
192.0.2.6/32         ldp   MPLS    -        9      192.168.67.1  1000
192.0.2.8/32         ldp   MPLS    -        9      192.168.57.1  2000
192.0.2.9/32         ldp   MPLS    -        9      192.168.79.2  1000
192.0.2.10/32        ldp   MPLS    -        9      192.168.107.2 1000
-----
Flags: B = BGP backup route available
       E = inactive best-external BGP route
=====
*A:PE-7#
```

Note that the LDP LSP metric follows the IGP cost.

- Enable LDP FRR on PE-7. This is a two-fold configuration command: first the IGP needs to be triggered to do LFA NH computation, and secondly, FRR needs to be enabled within the LDP context. Translated into configuration commands, this becomes:

```
*A:PE-7# configure router ospf loopfree-alternate
*A:PE-7# show router ospf status | match LFA
LFA                               : Enabled

*A:PE-7# configure router ldp fast-reroute

*A:PE-7# show router ldp status | match FRR
FRR                               : Enabled          Mcast Upstream FRR : Disabled
```

After issuing these two CLI commands, the software pre-computes both a primary and a backup Next-hop Label Forwarding Entry (NHLFE) for each LDP FEC in the network and downloads it to the IOM/IMM. The primary NHLFE corresponds to the label of the FEC received from the primary NH as per standard LDP resolution of the FEC prefix in the Routing Table Manager (RTM). The backup NHLFE corresponds to the label received for the same FEC from an LFA NH.

Configuration

The **show router route-table alternative** command adds an LFA flag to the associated alternative NH for a specific destination prefix. Other useful IGP related show commands are **show router ospf lfa-coverage** and **show router ospf routes alternative detail**.

```
*A:PE-7# show router route-table alternative
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                       Metric
  Alt-NextHop                                     Alt-
                                                Metric
-----
192.0.2.5/32                                     Remote OSPF    16h27m07s    10
    192.168.57.1                                 1000
    192.168.67.1 (LFA)                          2000
192.0.2.6/32                                     Remote OSPF    16h27m07s    10
    192.168.67.1                                 1000
    192.168.57.1 (LFA)                          2000
192.0.2.7/32                                     Local  Local   16h37m53s    0
    system                                       0
192.0.2.8/32                                     Remote OSPF    16h24m34s    10
    192.168.57.1                                 2000
    192.168.67.1 (LFA)                          2000
192.0.2.9/32                                     Remote OSPF    16h22m15s    10
    192.168.79.2                                 1000
    192.168.107.2 (LFA)                        2000
192.0.2.10/32                                    Remote OSPF    16h20m19s    10
    192.168.107.2                               1000
    192.168.79.2 (LFA)                        2000
192.168.56.0/30                                  Remote OSPF    16h27m07s    10
    192.168.57.1                                 2000
    192.168.67.1 (LFA)                        3000
192.168.57.0/30                                  Local  Local   16h29m11s    0
    int-PE-7-PE-5                               0
192.168.58.0/30                                  Remote OSPF    16h27m07s    10
    192.168.57.1                                 2000
    192.168.67.1 (LFA)                        3000
192.168.67.0/30                                  Local  Local   16h28m55s    0
    int-PE-7-PE-6                               0
192.168.68.0/30                                  Remote OSPF    16h27m07s    10
    192.168.67.1                                 2000
    192.168.57.1 (LFA)                        3000
192.168.78.0/30                                  Remote OSPF    16h24m27s    10
    192.168.57.1                                 3000
    192.168.67.1 (LFA)                        3000
192.168.79.0/30                                  Local  Local   16h28m16s    0
    int-PE-7-PE-9                               0
192.168.89.0/30                                  Remote OSPF    16h22m10s    10
    192.168.79.2                                 2000
    192.168.107.2 (LFA)                      3000
192.168.107.0/30                                 Local  Local   16h27m52s    0
    int-PE-7-PE-10                              0
192.168.109.0/30                                 Remote OSPF    16h22m10s    10
    192.168.79.2                                 2000
    192.168.107.2 (LFA)                      3000
-----
*A:PE-7#
```

Displaying the Label Forwarding Information Base (LFIB) on PE-7 shows the available alternate NHs; displayed with BU flag.

```
*A:PE-7# show router ldp bindings active
=====
Legend:  (S) - Static          (M) - Multi-homed Secondary Support
         (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op   IngLbl   EgrLbl   EgrIntf/LspId  EgrNextHop
-----
192.0.2.5/32          Push  --       262143   1/1/4           192.168.57.1
192.0.2.5/32          Push  --       262142BU 1/1/3           192.168.67.1
192.0.2.5/32          Swap 262142   262143   1/1/4           192.168.57.1
192.0.2.5/32          Swap 262142   262142BU 1/1/3           192.168.67.1
192.0.2.6/32          Push  --       262143   1/1/3           192.168.67.1
192.0.2.6/32          Push  --       262142BU 1/1/4           192.168.57.1
192.0.2.6/32          Swap 262141   262143   1/1/3           192.168.67.1
192.0.2.6/32          Swap 262141   262142BU 1/1/4           192.168.57.1
192.0.2.7/32          Pop   262143   --        --              --
192.0.2.8/32          Push  --       262140   1/1/4           192.168.57.1
192.0.2.8/32          Push  --       262140BU 1/1/3           192.168.67.1
192.0.2.8/32          Swap 262140   262140   1/1/4           192.168.57.1
192.0.2.8/32          Swap 262140   262140BU 1/1/3           192.168.67.1
192.0.2.9/32          Push  --       262143   1/1/1           192.168.79.2
192.0.2.9/32          Push  --       262138BU 1/1/2           192.168.107.2
192.0.2.9/32          Swap 262139   262143   1/1/1           192.168.79.2
192.0.2.9/32          Swap 262139   262138BU 1/1/2           192.168.107.2
192.0.2.10/32         Push  --       262143   1/1/2           192.168.107.2
192.0.2.10/32        Push  --       262138BU 1/1/1           192.168.79.2
192.0.2.10/32        Swap 262138   262143   1/1/2           192.168.107.2
192.0.2.10/32        Swap 262138   262138BU 1/1/1           192.168.79.2
=====
*A:PE-7#
```

Finally, a synchronization timer is enabled between the IGP and LDP protocol when LDP FRR is enabled. From the moment that the interface for the previous primary NH is restored, the IGP may re-converge back to that interface before LDP has completed the FEC exchange with its neighbor over that interface. This may cause LDP to de-program the LFA NH from the FEC and blackhole the traffic. In this example a timer of 10 seconds is used. Translated into configuration commands, this becomes:

```
*A:PE-x# configure router interface <itf-name> ldp-sync-timer 10
```

When this timer is set, when a failed interface is subsequently restored, the IGP advertises this link into the network with an infinite metric for the period of this timer. When the failed link is restored, the **ldp-sync-timer** is started, and LDP adjacencies are brought up over the restored link and a label exchange is completed between the peers. After the **ldp-sync-timer expires**, the normal metric is advertised into the network again.

At this point, everything is in place to start creating LFA policies to influence the calculated LFA NHs.

Step 2. Create a route-next-hop policy template.

This is a mandatory step in the context of LFA policies. The route-next-hop template name is maximum of 32 characters long. Creating a route-next-hop policy is done in the following way:

```
*A:PE-x# configure router route-next-hop-policy template <template name>
```

Commands within a **route-next-hop** policy template follow the **begin-abort-commit** model. After a **commit**, the IGP re-evaluates the template and schedules a new LFA SPF to re-compute the LFA NH for the prefixes associated with this template.

Step 3. Configure admin-group constraints in route-next-hop policy.

This is an optional step in the context of LFA policies. Firstly, configure a group-name and a group-value, of each admin-group locally on the router. Translated into configuration commands:

```
*A:PE-x# configure router if-attribute admin-group <group-name> value <group-value>
```

Second, configure the admin-group membership of the IP interface(s) (network, IES or VPRN). Up to five admin-groups can be applied to an IP interface in one command but the command can be applied multiple times. The configured IP admin-group membership is applied in all levels/ areas the interface is participating in. Translated into configuration commands:

```
*A:PE-x# configure router interface <itf-name> if-attribute admin-group <group-name> [  
<group-name> ... (upto 5 max)]  
*A:PE-x# configure service vprn <svc-id> interface <itf-name> if-attribute admin-group  
<group-name> [ <group-name> ... (upto 5 max)]  
*A:PE-x# configure service ies <svc-id> interface <itf-name> if-attribute admin-group  
<group-name> [ <group-name> ... (upto 5 max)]
```

Third, add the IP admin-group constraints into the route-next-hop policy template one by one. The **include-group** statement instructs the LFA SPF selection algorithm to select a subset of LFA NHs among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. The **pref** option is used to provide a relative preference for the admin group selection. A lower preference value means that LFA SPF will first attempt to select an LFA backup NH which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured, then it is the least preferred (default preference value is 255).

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member.

The **exclude-group** command simply prunes all links belonging to the specified admin group before making the LFA backup NH selection for a prefix. If the same group name is part of both **include** and **exclude** statements, the exclude statement will take precedence. In other words, the **exclude** statement can be viewed as having an implicit preference value of 0.

Translated into configuration commands, this becomes:

```
*A:PE-x# configure router route-next-hop-policy template <template-name> exclude-group
<group-name>
*A:PE-x# configure router route-next-hop-policy template <template-name> include-group
<group-name> [pref <preference>]
```

Step 4. Configure SRLG constraints in route-next-hop policy.

This is an optional step in the context of LFA policies. Firstly, configure a group-name and group-value, of each SRLG group locally on the router. Translated into configuration commands this becomes:

```
*A:PE-x# configure router if-attribute srlg-group <group-name> value <group-value>
```

Second, configure the SRLG group membership of the IP interfaces (network, IES or VPRN). Up to five SRLG groups can be applied to an IP interface in one command but the command can be applied multiple times. The configured IP SRLG group membership is applied in all levels/areas the interface is participating in. Translated into configuration commands this becomes:

```
*A:PE-x# configure router interface <itf-name> if-attribute srlg-group <group-name> [
<group-name> ... (upto 5 max)]
*A:PE-x# configure service vprn <svc-id> interface <itf-name> if-attribute srlg-group
<group-name> [ <group-name> ... (upto 5 max)]
*A:PE-x# configure service ies <svc-id> interface <itf-name> if-attribute srlg-group
<group-name> [ <group-name> ... (upto 5 max)]
```

Third, add the IP SRLG group constraints into the route-next-hop policy template. When this command is applied to a prefix, the LFA SPF attempts to select an LFA NH which uses an outgoing interface that does not participate in any of the SRLGs of the outgoing interface used by the primary NH. Translated into configuration commands, this becomes:

```
*A:PE-x# configure router route-next-hop-policy template <template-name> srlg-enable
```

Step 5. Configure the protection type in route-next-hop policy.

This is an optional step in the context of LFA policies. With the use of LFA policies, the user can also select if link protection or node protection is preferred for IP prefixes and LDP FEC prefixes protected by a backup LFA NH. By default, node protection is chosen. The implementation falls back to link protection if no LFA NH is found for node protection. Translated into configuration commands, this becomes:

```
*A:PE-x# configure router route-next-hop-policy template <template-name> protection-type {link|node}
```

Step 6. Configure the NH preference type in route-next-hop policy.

This is an optional step in the context of LFA policies. With the use of LFA policies, the user can also select if tunnel backup NH or IP backup NH is preferred for IP prefixes and LDP FEC prefixes protected by a backup LFA NH. By default, IP backup NH is chosen. The implementation falls back to the other type (tunnel) if no LFA NH of the preferred type is found. Translated into configuration commands, this becomes:

```
*A:PE-x# configure router route-next-hop-policy template <template-name> nh-type {ip|tunnel}
```

Step 7. Apply the route-next-hop policy template to an IP interface.

When the route-next-hop policy is applied to an IP interface, all prefixes using this interface as primary NH take the selection criteria specified in Step 3, Step 4, Step 5 and Step 6 into account. Translated into configuration commands, this becomes:

```
*A:PE-x# configure router ospf area interface lfa-policy-map route-nh-template <template-name>
*A:PE-x# configure router ospf3 area interface lfa-policy-map route-nh-template <template-name>
*A:PE-x# configure service vprn ospf area interface lfa-policy-map route-nh-template <template-name>
*A:PE-x# configure service vprn ospf3 area interface lfa-policy-map route-nh-template <template-name>
```

Examples

All of the examples focus on providing another LFA NH for LDP FEC prefix 192.0.2.6/32 and 192.0.2.5/32 (the system IP addresses of PE-6 and PE-5) , with PE-7 being the Point of Local Repair (PLR).

See [Figure 279 on page 1788](#) for the network topology.

As shown earlier, the default LFA NH (without policy) for both LDP FEC prefixes is as follows:

```
*A:PE-7# show router ldp bindings active
=====
Legend:  (S) - Static          (M) - Multi-homed Secondary Support
         (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
```

Prefix	Op	IngLbl	EgrLbl	EgrIntf/LspId	EgrNextHop
192.0.2.5/32	Push	--	262143	1/1/4	192.168.57.1
192.0.2.5/32	Push	--	262142BU	1/1/3	192.168.67.1
192.0.2.5/32	Swap	262136	262143	1/1/4	192.168.57.1
192.0.2.5/32	Swap	262136	262142BU	1/1/3	192.168.67.1
192.0.2.6/32	Push	--	262143	1/1/3	192.168.67.1
192.0.2.6/32	Push	--	262142BU	1/1/4	192.168.57.1
192.0.2.6/32	Swap	262135	262143	1/1/3	192.168.67.1
192.0.2.6/32	Swap	262135	262142BU	1/1/4	192.168.57.1

This default LFA NH can be changed by adding specific selection criteria inside a route-next-hop policy template.

Examples

Example 1: admin-group

The objective is to force the LFA NH for both LDP FEC prefixes to use the path between PE-7 and PE-8.

Define admin-group 'red' with value '1' and apply it on the IP interfaces PE-7 to PE-5 and PE-7 to PE-6.

```
*A:PE-7# configure router if-attribute admin-group "red" value 1

*A:PE-7# configure router interface "int-PE-7-PE-5" if-attribute admin-group "red"
*A:PE-7# configure router interface "int-PE-7-PE-6" if-attribute admin-group "red"
```

Define a route-next-hop policy template 'example1', which excludes IP admin-group 'red'.

```
*A:PE-7# configure router route-next-hop-policy
*A:PE-7>config>router>route-nh# info
    begin
    template "example1"
        exclude-group "red"
    exit
    commit
```

From the moment that route-next-hop policy template 'example1' is applied to the OSPF interfaces towards PE-5 and PE-6, the LFA NHs for both LDP FEC prefixes change. They now both point to the PE-7 to PE-8 IP interface as LFA backup NH:

```
*A:PE-7# configure router ospf area 0 interface "int-PE-7-PE-5" lfa-policy-map route-nh-
template "example1"
*A:PE-7# configure router ospf area 0 interface "int-PE-7-PE-6" lfa-policy-map route-nh-
template "example1"
```

```
*A:PE-7# show router ldp bindings active
```

```
=====
Legend:  (S) - Static          (M) - Multi-homed Secondary Support
         (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
```

```
LDP Prefix Bindings (Active)
```

```
=====
Prefix                Op   IngLbl   EgrLbl   EgrIntf/LspId  EgrNextHop
-----
192.0.2.5/32          Push  --       262143   1/1/4           192.168.57.1
192.0.2.5/32          Push  --       262142BU 1/1/5           192.168.78.2
192.0.2.5/32          Swap 262136   262143   1/1/4           192.168.57.1
192.0.2.5/32          Swap 262136   262142BU 1/1/5           192.168.78.2
192.0.2.6/32          Push  --       262143   1/1/3           192.168.67.1
192.0.2.6/32          Push  --       262141BU 1/1/5           192.168.78.2
192.0.2.6/32          Swap 262135   262143   1/1/3           192.168.67.1
192.0.2.6/32          Swap 262135   262141BU 1/1/5           192.168.78.2
=====
```

Example 2: SRLG

The objective is to force the LFA NH for both LDP FEC prefixes to use the PE-7 to PE-8 path.

Define SRLG group 'blue' with value '2' and apply it to the IP interfaces PE-7 to PE-5 and PE-7 to PE-6.

```
*A:PE-7# configure router if-attribute srlg-group "blue" value 2
*A:PE-7# configure router interface "int-PE-7-PE-5" if-attribute srlg-group "blue"
*A:PE-7# configure router interface "int-PE-7-PE-6" if-attribute srlg-group "blue"
```

Define a route-next-hop policy template 'example2', where SRLG is enabled

```
*A:PE-7# configure router route-next-hop-policy
*A:PE-7>config>router>route-nh# info
    begin
    template "example2"
        srlg-enable
    exit
    commit
```

From the moment that route-next-hop policy template 'example2' is applied to the OSPF interfaces towards PE-5 and PE-6, the LFA NHs for both LDP FEC prefixes change. They will both point now to the PE-7 to PE-8 interface as LFA backup NH:

```
*A:PE-7# configure router ospf area 0 interface "int-PE-7-PE-5" lfa-policy-map route-nh-template "example2"
*A:PE-7# configure router ospf area 0 interface "int-PE-7-PE-6" lfa-policy-map route-nh-template "example2"
```

```
*A:PE-7# show router ldp bindings active
```

```
=====
Legend:  (S) - Static          (M) - Multi-homed Secondary Support
         (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
```

Prefix	Op	IngLbl	EgrLbl	EgrIntf/LspId	EgrNextHop
192.0.2.5/32	Push	--	262143	1/1/4	192.168.57.1
192.0.2.5/32	Push	--	262142BU	1/1/5	192.168.78.2
192.0.2.5/32	Swap	262136	262143	1/1/4	192.168.57.1
192.0.2.5/32	Swap	262136	262142BU	1/1/5	192.168.78.2
192.0.2.6/32	Push	--	262143	1/1/3	192.168.67.1
192.0.2.6/32	Push	--	262141BU	1/1/5	192.168.78.2
192.0.2.6/32	Swap	262135	262143	1/1/3	192.168.67.1
192.0.2.6/32	Swap	262135	262141BU	1/1/5	192.168.78.2

Example 3: NH-type

The objective is to force the LFA NH for IP prefix 192.0.2.6/32 to use an RSVP tunnel.

Enable IP FRR and setup an RSVP LSP tunnel¹ towards 192.0.2.6 with a strict MPLS path going over PE-7 to PE-9 to PE-8 to PE-6.

```
*A:PE-7# configure router ip-fast-reroute

*A:PE-7# configure router mpls
  interface "system"
    no shutdown
  exit
  interface "int-PE-7-PE-9"
    no shutdown
  exit
  path "P-PE-7-PE-9-PE-8-PE6"
    hop 10 192.168.79.2 strict
    hop 20 192.168.89.1 strict
    hop 30 192.168.68.1 strict
    no shutdown
  exit
  lsp "LSP-PE-7-PE-6"
    to 192.0.2.6
    primary "P-PE-7-PE-9-PE-8-PE6"
    exit
    no shutdown
  exit
  no shutdown
```

Enable RSVP shortcut within the IGP on PE-7 and indicate that the newly created RSVP LSP is a possible shortcut candidate for LFA backup NH only.

```
*A:PE-7# configure router ospf rsvp-shortcut

*A:PE-7# configure router mpls lsp "LSP-PE-7-PE-6" igp-shortcut lfa-only
```

Displaying the tunnel-table of PE-7 shows that an LDP LSP and an RSVP LSP is available towards PE-6:

```
*A:PE-7# show router tunnel-table 192.0.2.6
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId  Pref    Nexthop      Metric
-----
192.0.2.6/32         rsvp  MPLS   1         7       192.168.79.2 16777215
192.0.2.6/32         ldp   MPLS   -         9       192.168.67.1 1000
```

1. Since an RSVP LSP is setup between PE-7 and PE-6, MPLS/RSVP protocols also need to be enabled on all the corresponding IP interfaces along the MPLS path.

```
*A:PE-7# show router mpls lsp
=====
MPLS LSPs (Originating)
=====
LSP Name                               To           Tun      Fastfail  Adm  Opr
                               Id           Id       Config
-----
LSP-PE-7-PE-6                         192.0.2.6    1        No        Up   Up
-----
LSPs : 1
=====
*A:PE-7#
```

```
*A:PE-7# show router route-table alternative 192.0.2.6/32
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                    Type  Proto  Age          Pref
      Next Hop[Interface Name]                Metric
      Alt-NextHop                              Alt-
                                                Metric
-----
192.0.2.6/32                            Remote OSPF   00h02m44s  10
      192.168.67.1                          1000
      192.168.57.1 (LFA)                     2000
-----
*A:PE-7#
```

Define a route-next-hop policy template **example3**, where nh-type is set to **tunnel**.

```
*A:PE-7# configure router route-next-hop-policy
*A:PE-7>config>router>route-nh# info
      begin
      template "example3"
      nh-type tunnel
      exit
      commit
```

From the moment that route-next-hop policy template **example3** is applied to the OSPF interface towards PE-6, the LFA NH uses the RSVP tunnel. Note that the reference to the RSVP tunnel-ID (1) in the following show output corresponds with the tunnel-ID shown in the previous **show router tunnel-table 192.0.2.6** output:

```
*A:PE-7# configure router ospf area 0 interface "int-PE-7-PE-6" lfa-policy-map route-nh-
template "example3"

*A:PE-7# show router route-table alternative 192.0.2.6/32
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                    Type  Proto  Age          Pref
      Next Hop[Interface Name]                Metric
      Alt-NextHop                              Alt-
                                                Metric
-----
```

Examples

```
192.0.2.6/32                               Remote OSPF      00h01m49s 10
      192.168.67.1                          1000
      192.0.2.6 (LFA) (tunneled:RSVP:1)    65535
```

*A:PE-7#

*A:PE-7# show router fib 1 nh-table-usage

=====

FIB Next-Hop Summary

```
=====
```

IPv4/IPv6	Active	Available
IP Next-Hop	10	16383
Tunnel Next-Hop	1	993279

```
=====
```

*A:PE-7#

Example 4: Exclude Prefix

The objective is to force no LFA NH for LDP FEC prefix 192.0.2.5/32 where PE-7 is the PLR.

From the introduction of IP/LDP FRR implementation in SR-OS, it is possible to exclude an IGP interface, IGP area (OSPF) or IGP level (IS-IS) from the LFA SPF computation. The user also has the ability to exclude specific prefixes from the LFA SPF by using well-known prefix-lists and policy statements.

Translated into configuration commands, this becomes:

```
A:PE-7# configure router policy-options
      begin
      prefix-list "lo0-PE-5"
        prefix 192.0.2.5/32 exact
      exit
      policy-statement "PE-5-exclude-LFA"
        entry 10
          from
            prefix-list "lo0-PE-5"
          exit
          action accept
          exit
        exit
      exit
    exit
  commit
```

The configured policy statement is applied to the IGP protocol. From the moment that it is applied, the existing LFA NH entries for LDP FEC prefix 192.0.2.5/32 disappear instantly (compare with Example 1 above):

```
*A:PE-7# configure router ospf loopfree-alternate-exclude prefix-policy "PE-5-exclude-LFA"

*A:PE-7# show router ldp bindings active prefix 192.0.2.5/32
=====
Legend:  (S) - Static          (M) - Multi-homed Secondary Support
         (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op    IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----
192.0.2.5/32          Push  --        262143    1/1/4           192.168.57.1
192.0.2.5/32          Swap  262136    262143    1/1/4           192.168.57.1
-----
No. of Prefix Active Bindings: 2
=====
```

Conclusion

In production MPLS networks where IP FRR and/or LDP FRR is deployed it is possible that the existing calculated LFA NHs are not always taking the most optimal or desirable paths.

With LFA policies, operators have better control on the way in which LFA backup NHs are computed.

Different selection criteria can be part of the route-next-hop policy: IP admin-groups, IP SRLG groups, protection type preference and NH type preference.