

MPLS LDP FRR using ISIS as IGP

In This Chapter

This section describes MPLS LDP FRR using ISIS as the IGP.

Topics in this section include:

- [Applicability on page 1102](#)
- [Summary on page 1103](#)
- [Configuration on page 1104](#)
- [Conclusion on page 1124](#)

Applicability

MPLS Label Distribution Protocol Fast Re-Route (LDP FRR) is supported on all 7x50 platforms including the 7710 SR and 7710 SR c-4/12 or the 7750 SR c-4/12. This feature is supported on all IOM/IMMs and MDA/CMA types that support network interfaces from 9.0R4 and higher. This feature was tested on release 12.0.R5. There are no pre-requisites for this configuration.

Summary

LDP FRR improves convergence in case of a single link or single node failure in the network. Convergence times will be in the order of 10s of milliseconds. This is important to some application services (like VoIP) which are sensitive to traffic loss when running over the MPLS network.

Without FRR, link and/or node failures inside an MPLS LDP network result in traffic loss in the order of 100s of milliseconds. The reason for that is that LDP depends on the convergence of the underlying IGP (IS-IS sending LSPs in this case). After IGP convergence, LDP itself needs to compute new primary next-hop Label Forwarding Entries (NHLFEs) for all affected Forwarding Equivalence Classes (FECs). Finally, the different Label Forwarding Information Bases (LFIBs) are updated.

When FRR is configured on a node, the node pre-computes primary NHLFEs for all FECs and in addition it will pre-compute backup NHLFEs for all FECs. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop (see also RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*). Both primary NHLFEs and backup NHLFEs are programmed in the IOM/IMM which makes it possible to converge very quickly.

Implementation

The 7x50 software has implemented Inequality 1 (link criterion) and Inequality 3 (node criterion) of RFC 5286. Similar to the Shortest Path Tree (SPT) computation that is part of standard link-state routing functionality, also the LFA next-hop computation is based on the IGP metric.

The underlying LFA formulas appear in the following format:

Inequality 1: $[SP(\text{backup NHR}, D) < \{SP(\text{backup NHR}, S) + SP(S, D)\}]$

Inequality 3: $[SP(\text{backup NHR}, D) < \{SP(\text{backup NHR}, PN) + SP(PN, D)\}]$

With 'SP' = 'shortest IGP metric path', 'NHR' = 'next-hop router', 'D' = 'destination', 'S' = 'source node or upstream node doing the actual LFA next-hop computation' and 'PN' = 'protected node'. The inequality 3 rule is stricter than the inequality 1 rule. See [Additional Topics on page 1118](#) for a practical example on these formulas.

Configuration

This section provides information to configure:

- [Configuring the IP/MPLS network. on page 1104](#)
- [Enabling LDP FRR and verify with show commands. on page 1107](#)
- [Enable a synchronization timer between IGP and LDP protocol. on page 1112](#)
- [Data path verification using a Layer 2/VLL service. on page 1112](#)

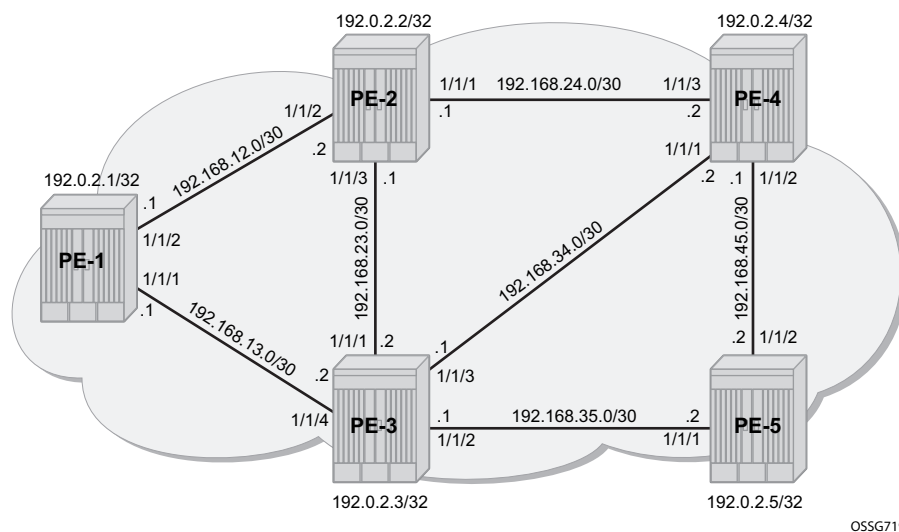


Figure 156: Initial Topology

Step 1. Configuring the IP/MPLS network.

The system addresses and IP interface addresses are configured according to [Figure 156](#). An interior gateway protocol (IGP) is needed to distribute routing information on all PEs. In our case, the IGP is IS-IS where each PE is acting as a Level 2 router. A configuration example is shown for PE-1. Similar configurations can be derived for the other PEs.

```
*A:PE-1# configure router isis
  level-capability level-2
  level 2
    wide-metrics-only
  exit
  interface "system"
  exit
  interface "int-PE-1-PE-2"
```

```

        interface-type point-to-point
    exit
    interface "int-PE-1-PE-3"
        interface-type point-to-point
    exit
    no shutdown

```

IS-IS interfaces are setup as type point-to-point to improve convergence since no DR/BDR election process is done. To verify that IS-IS adjacencies are up, **show router isis adjacency** is performed. To check if IP interface addresses/subnets are known on all PEs, **show router route-table** or **show router fib slot-number** will display the content of the forwarding information base (FIB).

```
*A:PE-1# show router isis adjacency
```

```

=====
ISIS Adjacency
=====
System ID                Usage State Hold Interface                MT-ID
-----
PE-2                      L2    Up    21  int-PE-1-PE-2                0
PE-3                      L2    Up    20  int-PE-1-PE-3                0
-----
Adjacencies : 2
=====

```

```
*A:PE-1#
```

```
*A:PE-1# show router route-table
```

```

=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type    Proto    Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.1/32                      Local   Local    07d23h34m    0
  system                               0
192.0.2.2/32                      Remote  ISIS     07d23h34m    18
  192.168.12.2                        10
192.0.2.3/32                      Remote  ISIS     07d23h17m    18
  192.168.13.2                        10
192.0.2.4/32                      Remote  ISIS     07d22h58m    18
  192.168.12.2                        20
192.0.2.5/32                      Remote  ISIS     07d23h17m    18
  192.168.13.2                        20
192.168.12.0/30                   Local   Local    07d23h34m    0
  int-PE-1-PE-2                       0
192.168.13.0/30                   Local   Local    07d23h34m    0
  int-PE-1-PE-3                       0
192.168.23.0/30                   Remote  ISIS     07d23h16m    18
  192.168.12.2                        20
192.168.24.0/30                   Remote  ISIS     07d23h34m    18
  192.168.12.2                        20
192.168.34.0/30                   Remote  ISIS     07d23h17m    18
  192.168.13.2                        20
192.168.35.0/30                   Remote  ISIS     07d23h17m    18

```

Configuration

```
192.168.13.2 20
192.168.45.0/30 Remote ISIS 03h59m10s 18
192.168.12.2 30
```

```
-----
No. of Routes: 12
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
```

```
*A:PE-1#
```

```
*A:PE-1# show router fib 1
```

```
=====
FIB Display
=====
```

Prefix [Flags]	NextHop	Protocol
192.0.2.1/32		LOCAL
192.0.2.1 (system)		
192.0.2.2/32	192.168.12.2 (int-PE-1-PE-2)	ISIS
192.0.2.3/32	192.168.13.2 (int-PE-1-PE-3)	ISIS
192.0.2.4/32	192.168.12.2 (int-PE-1-PE-2)	ISIS
192.0.2.5/32	192.168.13.2 (int-PE-1-PE-3)	ISIS
192.168.12.0/30	192.168.12.0 (int-PE-1-PE-2)	LOCAL
192.168.13.0/30	192.168.13.0 (int-PE-1-PE-3)	LOCAL
192.168.23.0/30	192.168.12.2 (int-PE-1-PE-2)	ISIS
192.168.24.0/30	192.168.12.2 (int-PE-1-PE-2)	ISIS
192.168.34.0/30	192.168.13.2 (int-PE-1-PE-3)	ISIS
192.168.35.0/30	192.168.13.2 (int-PE-1-PE-3)	ISIS
192.168.45.0/30	192.168.12.2 (int-PE-1-PE-2)	ISIS

```
-----
Total Entries : 12
=====
```

```
*A:PE-1#
```

Initially, the default IS-IS Level 2 metric is applied on all interfaces (value 10).

```
*A:PE-1# show router isis status | match "L2 Default Metric"
L2 Default Metric      : 10
*A:PE-1#
```

The next step in the process of setting up the IP/MPLS network is setting up interface-LDP sessions on all interfaces.

```
*A:PE-1# configure router ldp
      interface-parameters
        interface "int-PE-1-PE-2"
        exit
        interface "int-PE-1-PE-3"
        exit
      exit
      targeted-session
      exit
```

There is now a full mesh of LDP LSPs setup between all PE's system interfaces. As an example, the tunnel-table on PE-1 will look like this:

```
*A:PE-1# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId  Pref   Nexthop      Metric
-----
192.0.2.2/32         ldp  MPLS   -         9     192.168.12.2  10
192.0.2.3/32         ldp  MPLS   -         9     192.168.13.2  10
192.0.2.4/32         ldp  MPLS   -         9     192.168.12.2  20
192.0.2.5/32         ldp  MPLS   -         9     192.168.13.2  20
-----
Flags: B = BGP backup route available
       E = inactive best-external BGP route
=====
*A:PE-1#
```

Note that the LDP LSP metric follows the IGP cost. Optionally, LSP metrics can be applied but this is out of the scope for this configuration note.

Step 2. Enabling LDP FRR and verify with show commands.

Since LDP FRR is using LFA next-hop pre-computation by the IGP (as described in RFC 5286), the IGP CLI configuration will look like this:

```
*A:PE-1# configure router isis loopfree-alternate
*A:PE-1# show router isis status | match Loopfree
Loopfree-Alternate   : Enabled
*A:PE-1#
```

After enabling LFA inside the IGP context, FRR needs to be enabled within the LDP context:

```
*A:PE-1# configure router ldp fast-reroute
*A:PE-1# show router ldp status | match FRR
FRR                   : Enabled                Mcast Upstream FRR   : Disabled
*A:PE-1#
```

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

Note: For point-to-point interfaces, when multiple LFA next-hops are found for a given primary next-hop, the following selection algorithms are used:

- It will pick the node-protect type in favor of the link-protect type.
- If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
- If more than one LFA next-hop with the same cost, SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

Several show commands are possible to display LFA information:

- The **show router isis statistics** command displays the number of LFA runs on a specific node.

```
*A:PE-1# show router isis statistics
```

```
<snipped>
```

```
LFA Statistics
LFA Runs          : 25
```

```
<snipped>
```

```
*A:PE-1#
```

- The **show router isis lfa-coverage** command performs a mathematical calculation between the number of nodes and IPv4/IPv6 routes in the network versus present LFA next-hop protections. In our network (see [Figure 156](#)), all IS-IS links have a default Level 2 metric of 10. This results in all four nodes and all IS-IS routes learned by PE1 being 100% LFA protected (link or node). Refer to the following output.

```
*A:PE-1# show router isis lfa-coverage
```

```
=====
```

```
LFA Coverage
```

```
=====
```

Topology	Level	Node	IPv4	IPv6
IPv4 Unicast	L1	0/0 (0%)	9/9 (100%)	0/0 (0%)
IPv6 Unicast	L1	0/0 (0%)	0/0 (0%)	0/0 (0%)
IPv4 Multicast	L1	0/0 (0%)	0/0 (0%)	0/0 (0%)
IPv6 Multicast	L1	0/0 (0%)	0/0 (0%)	0/0 (0%)
IPv4 Unicast	L2	4/4 (100%)	9/9 (100%)	0/0 (0%)
IPv6 Unicast	L2	0/0 (0%)	0/0 (0%)	0/0 (0%)
IPv4 Multicast	L2	0/0 (0%)	0/0 (0%)	0/0 (0%)
IPv6 Multicast	L2	0/0 (0%)	0/0 (0%)	0/0 (0%)

```
=====
```


*A:PE-1#

- The **show router isis spf-log** command lists the SPF events indicating tye type (Lfa or Reg).

*A:PE-1# show router isis spf-log

=====
ISIS SPF Log
=====

When	Duration	L1 Nodes	L2 Nodes	Event Count	Type
09/23/2014 14:39:21	<0.01s	0	5	1	Reg
09/23/2014 14:39:22	<0.01s	-	-	1	Lfa
09/23/2014 14:40:01	<0.01s	0	5	1	Reg
09/23/2014 14:40:01	<0.01s	-	-	1	Lfa
09/23/2014 14:41:01	<0.01s	0	5	1	Reg
09/23/2014 14:41:01	<0.01s	-	-	1	Lfa
09/23/2014 14:43:17	<0.01s	0	5	1	Reg
09/23/2014 14:43:17	<0.01s	-	-	1	Lfa
09/23/2014 14:44:17	<0.01s	0	5	1	Reg
09/23/2014 14:44:17	<0.01s	-	-	1	Lfa
10/01/2014 09:32:35	<0.01s	0	5	1	Reg
10/01/2014 09:35:53	<0.01s	0	5	1	Reg
10/01/2014 09:35:53	<0.01s	-	-	1	Lfa
10/01/2014 09:43:05	<0.01s	0	5	1	Reg
10/01/2014 09:43:05	<0.01s	-	-	1	Lfa
10/01/2014 09:43:52	<0.01s	0	5	1	Reg
10/01/2014 09:43:52	<0.01s	-	-	1	Lfa
10/01/2014 13:41:50	<0.01s	0	5	1	Reg
10/01/2014 13:44:14	<0.01s	0	5	1	Reg
10/01/2014 13:44:14	<0.01s	-	-	1	Lfa

Log Entries : 20
=====

*A:PE-1#

- The **show router route-table** command adds an ‘L’ flag as reference that the associated prefix is having also an LFA next-hop available. For detailed interface address information used by the LFA calculation use the **show router route-table alternative** or **show router isis alternative** command. The output on PE-1 for PE-4 will look like this:

*A:PE-1# show router route-table 192.0.2.4

=====
Route Table (Router: Base)
=====

Dest Prefix[Flags] Next Hop[Interface Name]	Type	Proto	Age	Pref Metric
192.0.2.4/32 [L] 192.168.12.2	Remote	ISIS	07d23h01m	18 20

No. of Routes: 1

Flags: n = Number of times nexthop is repeated
B = BGP backup route available
L = LFA nexthop available
S = Sticky ECMP requested
=====

```

*A:PE-1#

*A:PE-1# show router route-table alternative 192.0.2.4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                               Type    Proto    Age          Pref
      Next Hop[Interface Name]                   Metric
      Alt-NextHop                                 Alt-
                                                Metric
-----
192.0.2.4/32                                       Remote  ISIS     07d23h02m  18
      192.168.12.2                                  20
      192.168.13.2 (LFA)                            20
-----

No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      Backup = BGP backup route
      LFA = Loop-Free Alternate nexthop
      S = Sticky ECMP requested
=====

*A:PE-1#

*A:PE-1# show router isis routes 192.0.2.4 alternative
=====
Route Table
=====
Prefix[Flags]                               Metric  Lvl/Typ    Ver.  SysID/Hostname
      NextHop                                MT      AdminTag
      Alt-Nexthop                             Alt-    Alt-Type
                                                Metric
-----
192.0.2.4/32                                       20      2/Int.    28    PE-2
      192.168.12.2                                  0       0
      192.168.13.2(L)                              20     NP
-----

No. of Routes: 1
Flags: L = Loop-Free Alternate nexthop
Legend: LP = linkProtection, NP = nodeProtection
=====

*A:PE-1#

```

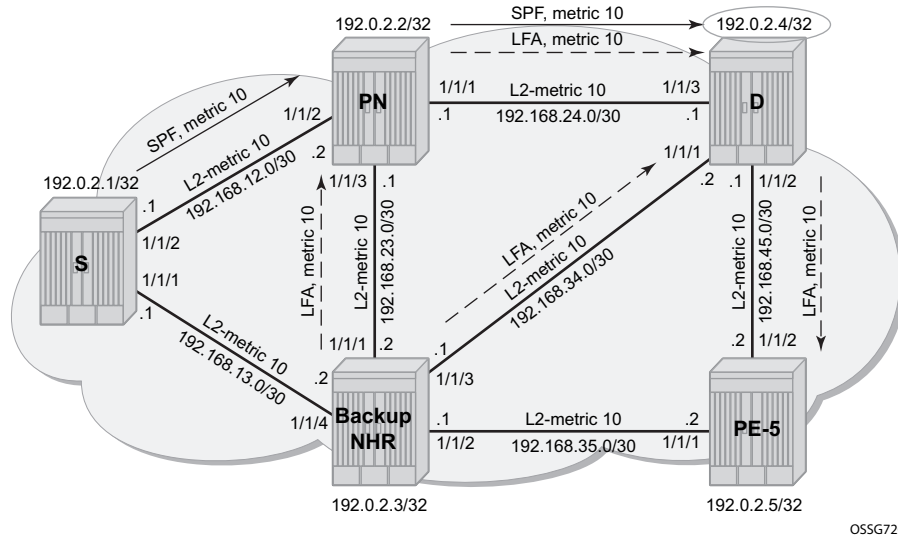
On PE-1, PE-4 (192.0.2.4/32) has a primary SPF next-hop pointing towards PE-2 (192.168.12.2) and an LFA next-hop pointing towards PE-3 (192.168.13.2).

Using the Inequality 3 formula on PE-1 for prefix 192.0.2.4/32, this becomes:

Inequality 3:

$$[SP(\text{backup NHR}, D) < \{SP(\text{backup NHR}, \text{PN}) + SP(\text{PN}, D)\}] \text{ or } [SP(\text{PE-3}, \text{PE-4}) < \{SP(\text{PE-3}, \text{PE-2}) + SP(\text{PE-2}, \text{PE-4})\}] \text{ or } [10 < \{10 + 10\}]$$

This means that Inequality 3 is met. The calculated LFA next-hop for prefix 192.0.2.4/32 on PE-1 is node-protecting PE-2, see Figure 81 for a graphical representation.



OSSG720

Figure 157: LFA Computation, Inequality 3 for Prefix PE-4 (D) on PE1 (S)

- The **show router ldp bindings** command displays the Label Information Base (LIB). A BU flag is present in case the associated label is used as backup NHLFE for the given prefix ¹. As an example, a display on PE-1 for prefix PE-4 will look like this:

```
*A:PE-1# show router ldp bindings prefix 192.0.2.4/32
=====
LDP LSR ID: 192.0.2.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Prefix Bindings
=====
Prefix          IngLbl      EgrLbl      EgrIntf/    EgrNextHop
Peer
-----
192.0.2.4/32    131068N    131068      1/1/2       192.168.12.2
  192.0.2.2
192.0.2.4/32    131068U    131068BU   1/1/1       192.168.13.2
  192.0.2.3
-----
No. of Prefix Bindings: 2
=====
*A:PE-1#
```

1. This is only possible because 7x50 LDP implementation is using liberal retention mode which means that every label mapping received by a peer is retained regardless of whether the LSR is the next-hop for the advertised mapping.

- The **show router ldp bindings active** command displays the Label Forwarding Information Base (LFIB). Also the BU flag is present and in addition a reference to the label action itself: **pop** for eLER, **push** for iLER and **swap** for LSR.

```
*A:PE-1# show router ldp bindings active prefix 192.0.2.4/32
=====
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
-----
192.0.2.4/32    Push  --        131068    1/1/2           192.168.12.2
192.0.2.4/32    Push  --        131068BU  1/1/1           192.168.13.2
192.0.2.4/32    Swap  131068    131068    1/1/2           192.168.12.2
192.0.2.4/32    Swap  131068    131068BU  1/1/1           192.168.13.2
-----
No. of Prefix Active Bindings: 4
=====
*A:PE-1#
```

Step 3. Enable a synchronization timer between IGP and LDP protocol.

Within an MPLS network using LDP it is common practice to enable a synchronization timer between LDP and the IGP. Also when LDP FRR is enabled, a situation can occur in which a synchronization timer between IGP and LDP will help: the revert scenario. When the interface for the previous primary next-hop is restored, IGP may re-converge before LDP completed the FEC exchange with its neighbor over that interface. This may cause LDP to de-program the LFA next-hop from the FEC and blackhole traffic.

In order to avoid these situations, it is recommended to first enable IGP-LDP synchronization on the LDP interface. The time is expressed in seconds and can have a value between 1 and 1800 seconds. Translated into configuration commands, this becomes:

```
*A:PE-1# configure router interface "int-PE-1-PE-2" ldp-sync-timer 10
*A:PE-1# configure router interface "int-PE-1-PE-3" ldp-sync-timer 10
```

When this timer is enabled, it means that when an interface is restored again, the IGP will advertise this link in the network with an infinite metric. The **ldp-sync-timer** is started, LDP adjacencies are brought up together with a label exchange. After the **ldp-sync-timer** expires, the normal metric is advertised in the network again.

Step 4. Data path verification using a Layer 2/VLL service.

Traffic generator ports are connected towards PE-1 and PE-5 for data verification, thus an Epipe service is created using an MPLS LDP based SDP on both PE-1 and PE-5.

```

*A:PE-1# configure service sdp 5
  far-end 192.0.2.5
  ldp
  keep-alive
  shutdown
  exit
  no shutdown

*A:PE-1# configure service epipe 1
  service-mtu 1450
  sap 1/1/3:1 create
  exit
  spoke-sdp 5:1 create
  no shutdown
  exit
  no shutdown
  
```

A similar configuration is configured on PE-5.

The IS-IS Level 2 metric value on the interface between PE-4 and PE-5 is decreased to 5, see Figure 158.

```

*A:PE-4# configure router isis interface int-PE-4-PE-5 level 2 metric 5
*A:PE-5# configure router isis interface int-PE-5-PE-4 level 2 metric 5
  
```

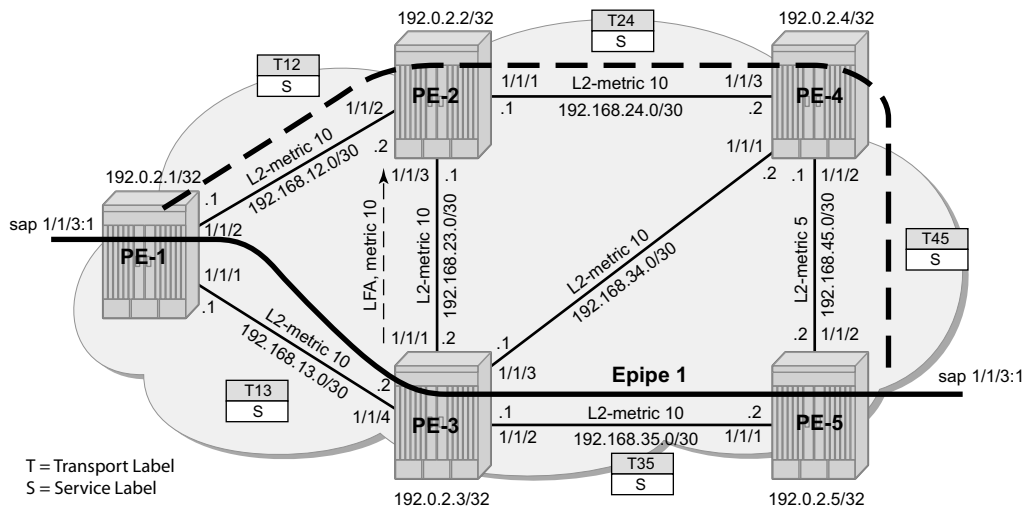


Figure 158: Data Verification, Direction PE-1 => PE-5 Using VLL Service

In this setup, PE-3 is node-protected for PE-5 prefix on PE-1:

```
*A:PE-1# show router isis routes alternative 192.0.2.5
=====
Route Table
=====
Prefix[Flags]                Metric    Lvl/Typ    Ver.  SysID/Hostname
  NextHop                    MT        AdminTag
Alt-Nexthop                  Alt-      Alt-Type
                               Metric
-----
192.0.2.5/32                  20        2/Int.     20    PE-3
  192.168.13.2                0         0
  192.168.12.2(L)             25        NP
-----
No. of Routes: 1
Flags: L = Loop-Free Alternate nexthop
Legend: LP = linkProtection, NP = nodeProtection
=====
*A:PE-1#
```

In normal conditions, MPLS traffic from PE-1 towards PE-5 over Epipe 1 will have two MPLS labels: 1) outer (transport) label given by LDP protocol, swapped on each intermediate LSR and 2) inner (service) label given by T-LDP, the same end-to-end. Refer to the following show commands.

The T-LDP service label is S (131066):

```
*A:PE-1# show router ldp bindings service-id 1
=====
LDP LSR ID: 192.0.2.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDPId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
E-Eth 1         1          5      192.0.2.5    131066U 131066S 1436 1436
-----
No. of VC Labels: 1
    <snipped>
*A:PE-1#
```

The transport LDP label between PE-1 and PE-3 for prefix 192.0.2.5/32 is T13 (131067):

```
*A:PE-1# 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  --        131067    1/1/1           192.168.13.2
192.0.2.5/32          Push  --        131067BU  1/1/2           192.168.12.2
192.0.2.5/32          Swap  131067    131067    1/1/1           192.168.13.2
192.0.2.5/32          Swap  131067    131067BU  1/1/2           192.168.12.2
-----
No. of Prefix Active Bindings: 4
=====
*A:PE-1#

```

The transport LDP label between PE-3 and PE-5 for prefix 192.0.2.5/32 is T35 (131071):

```

*A:PE-3# 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  --        131071    1/1/2           192.168.35.2
192.0.2.5/32          Swap  131067    131071    1/1/2           192.168.35.2
-----
No. of Prefix Active Bindings: 2
=====
*A:PE-3#

```

When PE-3 reboots, PE-1 performs an immediate swap to LFA next-hop for prefix 192.0.2.5/32 bypassing PE-3. The service label remains the same, only the transport labels can change on the network ports PE-1 <=> PE-2, PE-2 <=> PE-4 and PE-4 <=> PE-5². Refer to the following show commands.

The T-LDP service label is S (131066):

```

*A:PE-1# show router ldp bindings service-id 1
=====
LDP LSR ID: 192.0.2.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP Service FEC 128 Bindings

```

2. LDP FRR MPLS label stack will never contain more than two labels. This is different when compared to RSVP-TE FRR facility mode which uses a three-label MPLS stack.

Configuration

```
=====
Type   VCId      SvcId      SDPId      Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
E-Eth  1           1           5           192.0.2.5     131066U 131066S 1436 1436
-----
No. of VC Labels: 1
<snipped>
*A:PE-1#
```

The transport LDP label value between PE-1 and PE-2 for prefix 192.0.2.5/32 is the same label (previously tagged as BU) as before the node failure event: T12 (131067):

```
*A:PE-1# 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 --      131067  1/1/2           192.168.12.2
192.0.2.5/32    Swap 131067  131067  1/1/2           192.168.12.2
-----
No. of Prefix Active Bindings: 2
=====
*A:PE-1#
```

The transport LDP label between PE-2 and PE-4 for prefix 192.0.2.5/32 is T24 (131067):

```
*A:PE-2# 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 --      131067  1/1/1           192.168.24.2
192.0.2.5/32    Swap 131067  131067  1/1/1           192.168.24.2
-----
No. of Prefix Active Bindings: 2
=====
*A:PE-2#
```

The transport LDP label between PE-4 and PE-5 for prefix 192.0.2.5/32 is T45 (131071):

```
*A:PE-4# 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)
=====
```


MPLS LDP FRR using ISIS as IGP

Prefix	Op	IngLbl	EgrLbl	EgrIntf/LspId	EgrNextHop
192.0.2.5/32	Push	--	131071	1/1/2	192.168.45.2
192.0.2.5/32	Swap	131067	131071	1/1/2	192.168.45.2

No. of Prefix Active Bindings: 2

=====
*A:PE-4#

Additional Topics

Metric Change

Ensure the network is back to its initial topology with all the level 2 metrics back to 10 before applying the changes referenced below.

Suppose that the IS-IS Level 2 metric between PE-2 and PE-3 changes to 30, then 100% LFA coverage is no longer possible. Translated into configuration commands, this becomes:

```
*A:PE-3# configure router isis interface "int-PE-3-PE-2" level 2 metric 30
*A:PE-2# configure router isis interface "int-PE-2-PE-3" level 2 metric 30
```

On PE-1, Inequality 3 formula will find LFA next-hop coverages for prefix PE-4 and PE-5. Inequality formula 1 will find LFA next-hop coverages for prefix PE-4, PE-5 and the subnet between PE-4 and PE-5.

Both inequality formulas are visualized in [Figure 159](#) and [Figure 160](#) for prefix 192.0.2.5/32 (= PE-5) on PE-1 acting as the source node for LFA next-hop computation.




Figure 159: LFA Computation, Inequality 3 for Prefix PE-5 (D) on PE-1 (S)

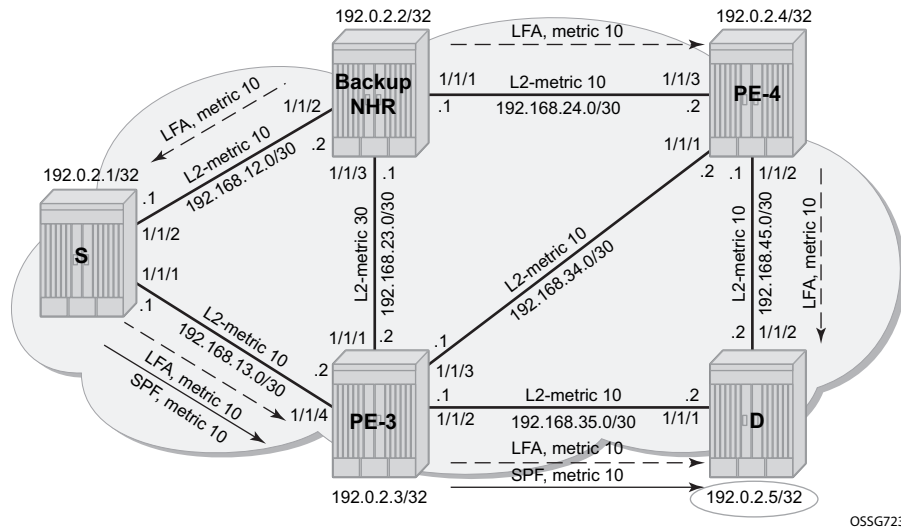


Figure 160: LFA Computation, Inequality 1 for Prefix PE-5 (D) on PE-1 (S)

Inequality 3 formula: $[SP(\text{backup NHR}, D) < \{SP(\text{backup NHR}, PN) + SP(PN, D)\}]$

For a node LFA next-hop calculation of prefix 192.0.2.5/32 (D) on PE-1, the above formula translated into text becomes:

[shortest path from backup next-hop router (PE-2) towards destination (PE-5) must be smaller than {shortest path from backup next-hop router (PE-2) towards protected node (PE-3) + shortest path from protected node (PE-3) to destination (PE-5)}].

The shortest path from backup next-hop router (PE-2) towards destination (PE-5) is going over PE-4, using IS-IS Level 2 metric 10 for interface PE-2 <=> PE-4 and IS-IS Level 2 metric 10 for interface PE-4 <=> PE-5. The shortest path from backup next-hop router (PE-2) towards protected node (PE-3) uses IS-IS Level 2 metric 30 for interface PE-2 <=> PE-3. The shortest path from protected node (PE-3) to destination (PE-5) uses IS-IS Level 2 metric 10 for interface PE-3 <=> PE-5. A concise example is displayed below.

```
Prefix 192.0.2.5/32 : SP(PE-2, PE-5) < SP(PE-2, PE-3) + SP(PE-3, PE-5)
                    10 + 10 < 30 + 10 => OK
```

Inequality 1 formula: $[SP(\text{backup NHR}, D) < \{SP(\text{backup NHR}, S) + SP(S, D)\}]$

For a link LFA next-hop calculation of prefix 192.0.2.5/32 (D) on PE-1, the formula translated displayed above into text becomes:

Additional Topics

[shortest path from backup next-hop router (PE-2) towards destination (PE-5) must be smaller than {shortest path from backup next-hop router (PE-2) towards source (PE-1) + shortest path from source (PE-1) to destination (PE-5)}].

The shortest path from backup next-hop router (PE-2) towards destination (PE-5) is going over PE-4, using IS-IS Level 2 metric 10 for interface PE-2 <=> PE-4 and IS-IS Level 2 metric 10 for interface PE-4 <=> PE-5. The shortest path from backup next-hop router (PE-2) towards source (PE-1) uses IS-IS Level 2 metric 10 for interface PE-2 <=> PE-1. The shortest path from source (PE-1) to destination (PE-5) follows the normal SPF calculation, going over PE-3, using IS-IS Level 2 metric 10 for interface PE-1 <=> PE-3 and IS-IS Level 2 metric 10 for interface PE-3 <=> PE-5. Written more concisely:

```
Prefix 192.0.2.5/32 : SP(PE-2,PE-5) < SP(PE-2,PE-1) + SP(PE-1,PE-5)
                    10 + 10      <    10 + (10 + 10)                => OK
```

For completion, all the other Inequality 1 calculations on PE-1 are given:

```
Prefix 192.0.2.2/32 : SP(PE-3,PE-2) < SP(PE-3,PE-1) + SP(PE-1,PE-2)
                    30      <    10 + 10                            => NOK
Prefix 192.0.2.3/32 : SP(PE-2,PE-3) < SP(PE-2,PE-1) + SP(PE-1,PE-3)
                    30      <    10 + 10                            => NOK
Prefix 192.0.2.4/32 : SP(PE-3,PE-4) < SP(PE-3,PE-1) + SP(PE-1,PE-2)
                    10      <    10 + 10                            => OK
Prefix 192.168.23.0/30 : SP(PE-3,D) < SP(PE-3,PE-1) + SP(PE-1,D)
                       30      <    10 + 10                            => NOK
Prefix 192.168.24.0/30 : SP(PE-3,D) < SP(PE-3,PE-1) + SP(PE-1,D)
                       30 + 10 <    10 + (10 + 10)                    => NOK
Prefix 192.168.34.0/30 : SP(PE-2,D) < SP(PE-2,PE-1) + SP(PE-1,D)
                       30 + 10 <    10 + (10 + 10)                    => NOK
Prefix 192.168.35.0/30 : SP(PE-2,D) < SP(PE-2,PE-1) + SP(PE-1,D)
                       30 + 10 <    10 + (10 + 10)                    => NOK
Prefix 192.168.45.0/30 : SP(PE-3,D) < SP(PE-3,PE-1) + SP(PE-1,D)
                       10 + 10 <    10 + (10 + 10 + 10)                => OK
```

As shown, only three are valid. On the 7x50, a summary command exists for LFA coverage on the router:

```
*A:PE-1# show router isis lfa-coverage
=====
LFA Coverage
=====
Topology          Level  Node          IPv4          IPv6
-----
IPV4 Unicast      L1    0/0 (0%)      3/9 (33%)    0/0 (0%)
IPV6 Unicast      L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Multicast    L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV6 Multicast    L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Unicast      L2    2/4 (50%)     3/9 (33%)    0/0 (0%)
IPV6 Unicast      L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Multicast    L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV6 Multicast    L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
=====
*A:PE-1#
```

IS-IS Overload Bit

Ensure the network is back to its initial topology with all the level 2 metrics back to 10 before applying the changes referenced below.

As stated in RFC 3137, *OSPF Stub Router Advertisement*, sometimes it is desirable not to have a router used as a transit node. For those cases, it is also desirable not to have that router used as transit node during the LFA next-hop computation. Within IS-IS protocol this is achieved by setting the overload bit. When other routers detect that this bit is set, they will only use this router for packets destined to the overloaded router's directly connected networks and IP prefixes.

As an example, setting of the IS-IS overload condition for a specific time on PE-2 provides following result on PE-1:

```
*A:PE-2# configure router isis overload timeout 60
## possible value 60 ..1800
## or
## infinity without a timeout value

*A:PE-1# show router isis lfa-coverage
=====
LFA Coverage
=====
Topology          Level  Node          IPv4          IPv6
-----
IPV4 Unicast      L1    0/0 (0%)      3/9 (33%)    0/0 (0%)
IPV6 Unicast      L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Multicast    L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV6 Multicast    L1    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Unicast      L2    1/4 (25%)     3/9 (33%)    0/0 (0%)
IPV6 Unicast      L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV4 Multicast    L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
IPV6 Multicast    L2    0/0 (0%)      0/0 (0%)    0/0 (0%)
=====

*A:PE-1#

*A:PE-1# 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
-----
192.0.2.1/32                Local  Local  07d23h49m 0
system                      0
192.0.2.2/32                Remote ISIS  07d23h48m 18
192.168.12.2                10
192.168.13.2 (LFA)         20
192.0.2.3/32                Remote ISIS  00h01m59s 18
192.168.13.2                10
192.0.2.4/32                Remote ISIS  00h00m40s 18
```

Additional Topics

```

192.168.13.2                                20
192.0.2.5/32                               Remote ISIS 00h01m59s 18
192.168.13.2                                20
192.168.12.0/30                             Local Local 07d23h48m 0
int-PE-1-PE-2                               0
192.168.13.0/30                             Local Local 07d23h48m 0
int-PE-1-PE-3                               0
192.168.23.0/30                             Remote ISIS 00h01m11s 18
192.168.12.2                                20
192.168.13.2 (LFA)                          30
192.168.24.0/30                             Remote ISIS 07d23h48m 18
192.168.12.2                                20
192.168.13.2 (LFA)                          30
192.168.34.0/30                             Remote ISIS 00h01m59s 18
192.168.13.2                                20
192.168.35.0/30                             Remote ISIS 00h01m59s 18
192.168.13.2                                20
192.168.45.0/30                             Remote ISIS 00h00m40s 18
192.168.13.2                                25

```

```

-----
No. of Routes: 12
Flags: n = Number of times nexthop is repeated
Backup = BGP backup route
LFA = Loop-Free Alternate nexthop
S = Sticky ECMP requested

```

```

=====
*A:PE-1#

```

```

*A:PE-1# show router isis routes alternative

```

```

-----
Route Table

```

```

=====
Prefix[Flags]                               Metric   Lvl/Typ   Ver.  SysID/Hostname
NextHop                                     MT       AdminTag
Alt-Nexthop                               Alt-     Alt-Type
                                           Metric
-----
192.0.2.1/32                                0        2/Int.    3     PE-1
0.0.0.0                                     0        0
192.0.2.2/32                                10       2/Int.    5     PE-2
192.168.12.2                               0        0
192.168.13.2(L)                            20       LP
192.0.2.3/32                                10       2/Int.    43    PE-3
192.168.13.2                               0        0
192.0.2.4/32                                20       2/Int.    47    PE-3
192.168.13.2                               0        0
192.0.2.5/32                                20       2/Int.    43    PE-3
192.168.13.2                               0        0
192.168.12.0/30                             10       2/Int.    4     PE-1
0.0.0.0                                     0        0
192.168.13.0/30                             10       2/Int.    43    PE-1
0.0.0.0                                     0        0
192.168.23.0/30                             20       2/Int.    45    PE-2
192.168.12.2                               0        0
192.168.13.2(L)                            30       LP
192.168.24.0/30                             20       2/Int.    5     PE-2
192.168.12.2                               0        0
192.168.13.2(L)                            30       LP

```

```

192.168.34.0/30          20          2/Int.      43    PE-3
   192.168.13.2          0          0
192.168.35.0/30        20          2/Int.      43    PE-3
   192.168.13.2          0          0
192.168.45.0/30        25          2/Int.      47    PE-3
   192.168.13.2          0          0
-----
No. of Routes: 12
Flags: L = Loop-Free Alternate nexthop
Legend: LP = linkProtection, NP = nodeProtection
=====
*A:PE-1#

```

On PE-1, only three Inequality 1 calculations are possible. Refer to the previous show commands.

```

[SP(backup NHR,D) < {SP(backup NHR,S) + SP(S,D)}]
      SP(PE-3,D) < SP(PE-3,PE-1) + SP(PE-1,D)
      10 + 10 < 10 + (10 + 10)                      => OK

```




Figure 161: IS-IS Overload on PE-2, Inequality 1 for 192.168.24.0/30 (D) on PE-1 (S)

Conclusion

In production MPLS networks where FRR needs to be deployed, a trade off must be made between RSVP-TE FRR versus LDP FRR. The two main advantages of using LDP FRR compared to RSVP FRR are simple configuration and LFA next-hop calculation is a local decision, which means no interoperability issues when working in a multi-vendor environment. The main disadvantage of using LDP FRR is that LFA next-hop calculation has to deal with source-route paradigm (inequality formulas exclude a path going over the original source router).