Inter-Area TE Point-to-Point LSPs

In This Chapter

This section describes inter-area TE point-to-point LSP configurations.

Topics in this section include:

Applicability

Inter-Area Traffic Engineering (TE) point-to-point (P2P) LSPs are supported on all 7x50 platforms, including the 7710 and 7750 SR-c4/12. This feature is supported on all IOM/IMM types. The configuration was tested on release 11.0.R4.

Summary

MPLS TE is implemented on a wide scale in current ISP networks to steer traffic across the backbone to facilitate efficient use of available bandwidth between the routers and to guarantee fast convergence in case a link or node fails.

Previously, the MPLS TE designs allowed for TE LSPs that are confined to only a single IGP area/level. This is due to the fact that the head-end has information in the TE database of only the local area (OSPF) or level (ISIS).

Inter-Area TE LSP Based On Explicit Route Expansion

To be able to support Inter Area MPLS traffic engineering, the design needs to be extended. Inter-Area TE LSP based on Explicit route Object (ERO) expansion allows for the head-end to calculate the ERO path within its own area/level and keep the remaining ABRs of other areas/levels as loose hops in the ERO path. On receiving a PATH message with a loose hop ERO, based on local configuration each ABR does a partial Constrained Shortest Path First (CSPF) calculation to the next ABR or full CSPF to reach the final destination.

Automatic selection of ABRs is supported, in this way the head-end node can work with an empty primary path. When the **to** field of an LSP definition is in a different area/level than the head-end node, CSPF will automatically compute the segment to the exit ABR router which advertised the prefix and which is currently the best path for resolving the prefix in Route Table Manager (RTM).

ABR Protection

Link and Node protection within the respective areas are supported through the TE capabilities of the IGP and RSVP in each area. To support ABR node protection, a bypass is required from the Point of Local Repair (PLR; node prior to ABR) to the Merge Point (MP; next-hop node to ABR). Two methods are possible: dynamic ABR protection and static ABR protection. Static ABR protection uses Manual Bypass Tunnels (MBTs), statically configured by the operator between PLR and MP.

For dynamic ABR protection, node ID propagation and signaling of an Exclude Route (XRO) object in RSVP PATH messages must both be supported.

Since the description of the RRO Node ID sub-object in RFC 4561 (*Definition of a Record Route Object (RRO) Node-Id Sub-Object*) is not clear about the format of the included node-address (S), interface-address (I) and label (L), the system is programmed to understand multiple formats: IL, SL, ISL, SIL, SLI, ILSL and SLIL. The system uses the SLIL (node-adddress, label, interface-address, label) format to include the node-ID itself.

XRO object inclusion (RFC 4874, *Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering*) in bypass RSVP PATH messages is required to exclude the protected ABR from the bypass path. The XRO object is filled in with ABRs system IP address.

Overview

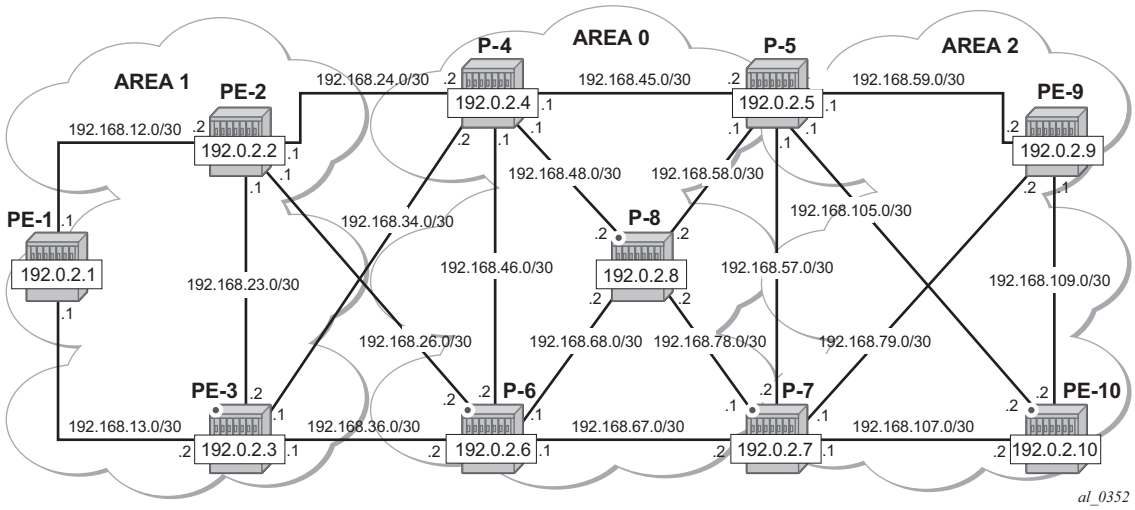

Figure 147: Inter-Area TE LSP Setup

The setup in this section contains 10 nodes in three areas. Figure 147 shows the physical topology of the setup.

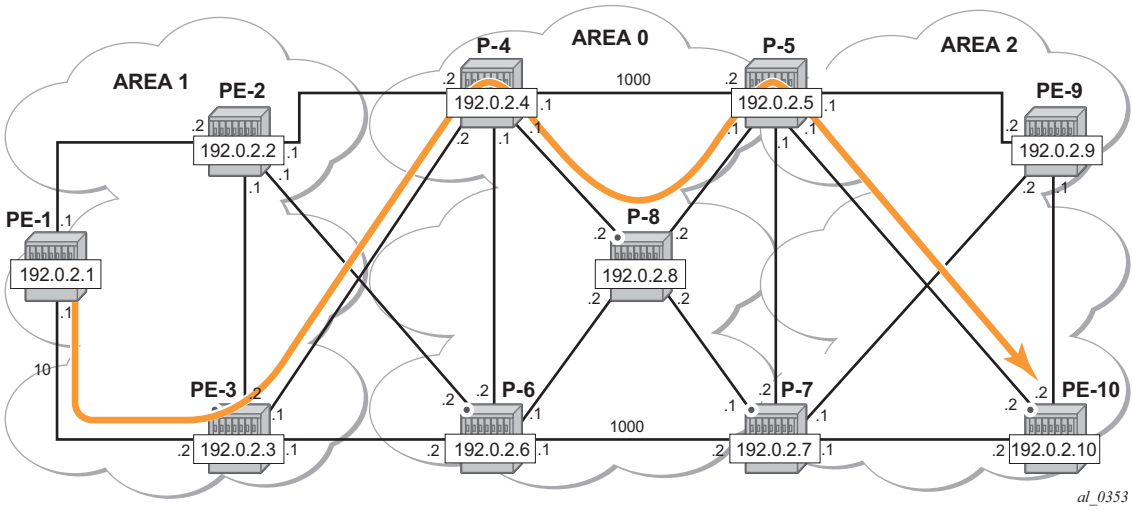

Figure 148: Inter-Area TE LSP Path

Figure 148 shows the LSP path intended to be setup through the network. An empty MPLS path is used. At the head-end node (PE-1), the destination address (PE-10) is learned via ABR node P-4 and ABR node P-5.

Configuration

The assumption is made that following base configuration has been implemented on the PEs:

- Cards, MDAs and ports configured
- Interfaces configured
- IGP areas configured and converged
- Traffic Engineering configured for the IGP
- MPLS and RSVP configured on all links in the network

OSPF or ISIS can be configured as the IGP; OSPF is used here.

```
*A:PE-1# show router ospf opaque-database
===============================================================================
OSPF Opaque Link State Database (Type : All)
===============================================================================
Type  Id              Link State Id     Adv Rtr Id      Age  Sequence   Cksum
-------------------------------------------------------------------------------
Area  0.0.0.1         1.0.0.1           192.0.2.1       357  0x800000ef 0xb622
Area  0.0.0.1         1.0.0.2           192.0.2.1       619  0x800001dc 0xc193
Area  0.0.0.1         1.0.0.3           192.0.2.1       461  0x800001dd 0xe42
Area  0.0.0.1         1.0.0.1           192.0.2.2       1434 0x800000f5 0xae22
Area  0.0.0.1         1.0.0.2           192.0.2.2       498  0x800001e8 0x85c3
Area  0.0.0.1         1.0.0.3           192.0.2.2       285  0x800001e3 0x92a2
Area  0.0.0.1         1.0.0.4           192.0.2.2       428  0x800001e7 0xd854
Area  0.0.0.1         1.0.0.5           192.0.2.2       713  0x800001e9 0x7ba8
Area  0.0.0.1         1.0.0.1           192.0.2.3       424  0x800000f1 0xba18
Area  0.0.0.1         1.0.0.2           192.0.2.3       763  0x800001e3 0xcb7f
Area  0.0.0.1         1.0.0.3           192.0.2.3       183  0x800001e5 0x6ac8
Area  0.0.0.1         1.0.0.4           192.0.2.3       294  0x800001e4 0x6fab
Area  0.0.0.1         1.0.0.5           192.0.2.3       763  0x800001ea 0xa04
Area  0.0.0.1         1.0.0.1           192.0.2.4       4    0x800000ed 0xc60e
Area  0.0.0.1         1.0.0.5           192.0.2.4       1395 0x800001e1 0x9a97
Area  0.0.0.1         1.0.0.6           192.0.2.4       1030 0x800001e1 0x3dde
Area  0.0.0.1         1.0.0.1           192.0.2.6       1367 0x800000ee 0xcc03
Area  0.0.0.1         1.0.0.5           192.0.2.6       535  0x800001e2 0xbd58
Area  0.0.0.1         1.0.0.6           192.0.2.6       709  0x800001de 0xf1f
-------------------------------------------------------------------------------
No. of Opaque LSAs: 19
===============================================================================
```

The output above shows the opaque database of PE-1. The information is only about routers that are part of area 0.0.0.1. PE-1 cannot calculate an end-to-end CSPF path to node PE-10 since this would require TE topology information from area 0.0.0.0 and area 0.0.0.2.

Each node announces its router-ID and each attached link that is part of that area, hence the 19 opaque LSAs in area 0.0.0.1.

Note in Figure 148 that the LSP should pass through node PE-3 and node P-8. In order to prefer a dynamic path from PE-1 to P-4 via PE-3 rather than through PE-2, it is necessary to configure on

PE-1 a lower IGP metric on the interface to PE-3 (the default metric is derived from the interface speed; in this case the metric is 100 by default).

```
*A:PE-1>config>router>ospf# area 1 interface "int-PE-1-PE-3" metric 10
```

Similarly, in the core, the IGP metric between P-4 <=> P-5 and P-6 <=> P-7 is increased to force the LSP to pass the core P-8 node.

```
*A:P-4>config>router>ospf# area 0 interface "int-P-4-P-5" metric 1000
*A:P-6>config>router>ospf# area 0 interface "int-P-6-P-7" metric 1000
```

MPLS Path Configuration

Since automatic ABR selection is performed, an empty MPLS path is enough on the head-end node PE-1. Using an empty MPLS path will ease the provisioning process and brings consistency since this empty MPLS path can be used for both intra and inter-area/level type LSPs.

```
*A:PE-1# configure router mpls
*A:PE-1>config>router>mpls# path "path-PE-10" no shutdown
```

MPLS LSP Configuration

Configure an LSP on PE-1 to PE-10 and include the previously created MPLS path as primary path. Enable CSPF and Fast Reroute (FRR) facility on the LSP.

```
*A:PE-1>config>router>mpls#
*A:PE-1>config>router>mpls# lsp "LSP-PE-1-PE-10" to 192.0.2.10
*A:PE-1>config>router>mpls# lsp "LSP-PE-1-PE-10" cspf
*A:PE-1>config>router>mpls# lsp "LSP-PE-1-PE-10" fast-reroute facility
*A:PE-1>config>router>mpls# lsp "LSP-PE-1-PE-10" primary "path-PE-10" no shutdown
*A:PE-1>config>router>mpls# lsp "LSP-PE-1-PE-10" no shutdown
```

At this stage the LSP is in an operational Down state with a failure code of noCspfRouteToDestination.

In order to get around the intra-area CSPF confinement, enable the ERO-expansion feature on all possible ABR nodes.

```
*A:P-4# configure router mpls cspf-on-loose-hop
*A:P-6# configure router mpls cspf-on-loose-hop

*A:P-7# configure router mpls cspf-on-loose-hop
*A:P-5# configure router mpls cspf-on-loose-hop
```

Note that cspf-on-loose-hop is only required if FRR or TE parameters are configured on the LSP. If any of these parameters is configured on the LSP and one of the ABRs along the path is not configured with cspf-on-loose-hop, the LSP will stay operationally down with Failure Code: badNode and an indication of the interface address of the Failure Node.

```
*A:PE-1# show router mpls lsp "LSP-PE-1-PE-10" path detail
...
From        : 192.0.2.1                         To             : 192.0.2.10
Adm State   : Up                                Oper State     : Down
…
Failure Code: badNode                           Failure Node: 192.168.24.2
```

The LSP path can also contain other strict and/or loose hops. Note however that cspf-on-loose-hop must be configured under MPLS whenever loose hops are configured in the MPLS path. This

command is needed to trigger ERO expansion and is only required for inter-area LSPs on all possible ABR nodes and all nodes not belonging to the 'ingress' area (namely, the same area as the iLER) which have a 'loose hop' reference in the MPLS path. However, for simplicity it can be configured on all nodes without having a negative effect.

The following trace shows the ERO calculation on the head-end to the first ABR.

```
*A:PE-1# debug router rsvp packet path detail
2 2013/08/20 16:29:20.89 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.1, To:192.0.2.10
          TTL:255, Checksum:0x88f3, Flags:0x0
Session    - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
SessAttr   - Name:LSP-PE-1-PE-10::path-PE-10
             SetupPri:7, HoldPri:0, Flags:0x17
RSVPHop    - Ctype:1, Addr:192.168.13.1, LIH:3
TimeValue  - RefreshPeriod:30
SendTempl  - Sender:192.0.2.1, LspId:56852
SendTSpec  - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
             MPU:20, MTU:1564
LabelReq   - IfType:General, L3ProtID:2048
RRO        - IpAddr:192.168.13.1, Flags:0x0
ERO        - IPv4Prefix 192.168.13.2/32, Strict
             IPv4Prefix 192.168.34.2/32, Strict
             IPv4Prefix 192.0.2.10/32, Loose
FRRObj     - SetupPri:7, HoldPri:0, HopLimit:16, BW:0.000 bps, Flags:0x2
             ExcAny:0x0, IncAny:0x0, IncAll:0x0
"
```

On the P-4 ABR the ERO is expanded to include the nodes of area 0.0.0.0 of which P-4 is also part. The RRO contains all the hops the PATH message has passed so far.

```
*A:P-4# debug router rsvp packet path detail
11 2013/08/20 08:54:40.97 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.1, To:192.0.2.10
          TTL:253, Checksum:0x176f, Flags:0x0
Session    - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
SessAttr   - Name:LSP-PE-1-PE-10::path-PE-10
             SetupPri:7, HoldPri:0, Flags:0x17
RSVPHop    - Ctype:1, Addr:192.168.48.1, LIH:4
TimeValue  - RefreshPeriod:30
SendTempl  - Sender:192.0.2.1, LspId:56852
SendTSpec  - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
             MPU:20, MTU:1564
LabelReq   - IfType:General, L3ProtID:2048
RRO        - IpAddr:192.168.48.1, Flags:0x0
             IpAddr:192.168.34.1, Flags:0x0
             IpAddr:192.168.13.1, Flags:0x0
ERO        - IPv4Prefix 192.168.48.2/32, Strict
             IPv4Prefix 192.168.58.1/32, Strict
             IPv4Prefix 192.0.2.10/32, Loose
FRRObj     - SetupPri:7, HoldPri:0, HopLimit:16, BW:0.000 bps, Flags:0x2
             ExcAny:0x0, IncAny:0x0, IncAll:0x0
"
```

Finally, the P-5 ABR will expand the ERO to the final destination PE-10:

```
*A:P-5# debug router rsvp packet path detail
4 2013/08/20 08:25:04.70 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.1, To:192.0.2.10
          TTL:251, Checksum:0xbfcd, Flags:0x0
Session    - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
SessAttr   - Name:LSP-PE-1-PE-10::path-PE-10
             SetupPri:7, HoldPri:0, Flags:0x17
RSVPHop    - Ctype:1, Addr:192.168.105.1, LIH:5
TimeValue  - RefreshPeriod:30
SendTempl  - Sender:192.0.2.1, LspId:56852
SendTSpec  - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
             MPU:20, MTU:1564
LabelReq   - IfType:General, L3ProtID:2048
RRO        - IpAddr:192.168.105.1, Flags:0x0
             IpAddr:192.168.58.2, Flags:0x0
             IpAddr:192.168.48.1, Flags:0x0
             IpAddr:192.168.34.1, Flags:0x0
             IpAddr:192.168.13.1, Flags:0x0
ERO        - IPv4Prefix 192.168.105.2/32, Strict
FRRObj     - SetupPri:7, HoldPri:0, HopLimit:16, BW:0.000 bps, Flags:0x2
             ExcAny:0x0, IncAny:0x0, IncAll:0x0
"
```

The MPLS LSP is now operational up and the LSP path can be shown in detail on the head-end, PE-1:

```
*A:PE-1# show router mpls lsp "LSP-PE-1-PE-10" path detail
===============================================================================
MPLS LSP LSP-PE-1-PE-10 Path  (Detail)
===============================================================================
Legend :
    @ - Detour Available              # - Detour In Use
    b - Bandwidth Protected           n - Node Protected
    s - Soft Preemption
    S - Strict                        L - Loose
    A - ABR
===============================================================================
-------------------------------------------------------------------------------
LSP LSP-PE-1-PE-10 Path path-PE-10
-------------------------------------------------------------------------------
LSP Name    : LSP-PE-1-PE-10                    Path LSP ID : 56852
From        : 192.0.2.1                         To          : 192.0.2.10
Adm State   : Up                                Oper State  : Up
Path Name   : path-PE-10                        Path Type   : Primary
Path Admin  : Up                                Path Oper   : Up
OutInterface: 1/1/1                             Out Label   : 131071
Path Up Time: 0d 01:05:21                       Path Dn Time: 0d 00:00:00
Retry Limit : 0                                 Retry Timer : 30 sec
RetryAttempt: 0                                 NextRetryIn : 0 sec

Adspec      : Disabled                          Oper Adspec : Disabled
CSPF        : Enabled                           Oper CSPF   : Enabled
Least Fill  : Disabled                          Oper LeastF*: Disabled
FRR         : Enabled                           Oper FRR    : Enabled
```

```
FRR NodePro*: Enabled                          Oper FRR NP : Enabled
FR Hop Limit: 16                               Oper FRHopL*: 16
FR Prop Adm*: Disabled                         Oper FRProp*: Disabled
Prop Adm Grp: Disabled                         Oper PropAG : Disabled
Inter-area  : True

Neg MTU     : 1560                             Oper MTU    : 1560
Bandwidth   : No Reservation                   Oper Bw     : 0 Mbps
Hop Limit   : 255                              Oper HopLim*: 255
Record Route: Record                           Oper RecRou*: Record
Record Label: Record                           Oper RecLab*: Record
SetupPriori*: 7                                Oper SetupP*: 7
Hold Priori*: 0                                Oper HoldPr*: 0
Class Type  : 0                                Oper CT     : 0
Backup CT   : None
MainCT Retry: n/a
    Rem     :
MainCT Retry: 0
    Limit   :
Include Grps:                                  Oper InclGr*:
None                                          None
Exclude Grps:                                  Oper ExclGr*:
None                                          None

Adaptive    : Enabled                          Oper Metric : 110
Preference  : n/a
Path Trans  : 21                               CSPF Queries: 11
Failure Code: noError                          Failure Node: n/a
ExplicitHops:
    No Hops Specified
Actual Hops :
    192.168.13.1 (192.0.2.1) @ n               Record Label        : N/A
 -> 192.168.13.2 (192.0.2.3) @ n               Record Label        : 131071
 -> 192.168.34.2 (192.0.2.4) @ n               Record Label        : 131071
 -> 192.0.2.8 (192.0.2.8) @ n                  Record Label        : 131070
 -> 192.168.48.2  @ n                          Record Label        : 131070
 -> 192.0.2.5 (192.0.2.5) @                    Record Label        : 131071
 -> 192.168.58.1  @                            Record Label        : 131071
 -> 192.0.2.10 (192.0.2.10)                    Record Label        : 131071
 -> 192.168.105.2                              Record Label        : 131071
ComputedHops:
    192.168.13.1(S)
 -> 192.168.13.2(S)
 -> 192.168.34.2(SA)
 -> 192.0.2.10(L)
ResigEligib*: False
LastResignal: n/a                              CSPF Metric : 110
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

ABR Node Protection

At this stage, the LSP is configured with facility FRR protection; link and node protection will be offered within each area. Dynamic ABR node protection requires the setup of a bypass tunnel from the PLR (node just upstream of the ABR) to the MP (node just downstream of the ABR). Two things are required for this:

Firstly, the PLR node (part of area x) needs to know the system IP address of MP node (part of area y) to setup the bypass. For this reason, the node-ID of the MP is included in the RESV message so that the PLR can link the manual bypass tunnel to the primary path to protect the ABR.

Secondly, the other ABR node receiving the RSVP bypass PATH message for the protected ABR needs to do an ERO expansion towards MP node. For this reason, the XRO object is included in the RSVP bypass PATH message, containing the node-ID of the protected ABR. As an example, a bypass PATH message is shown below on node PE-3.

The XRO object includes the system IP address of the protected ABR node (P-4) and the ERO object has MP node (P-8) as loose destination:

```
*A:PE-3# debug router rsvp packet path detail
13 2013/08/20 07:56:10.87 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.3, To:192.0.2.8
          TTL:17, Checksum:0x907f, Flags:0x0
Session    - EndPt:192.0.2.8, TunnId:61445, ExtTunnId:192.0.2.3
SessAttr   - Name:bypass-node192.0.2.4
             SetupPri:7, HoldPri:0, Flags:0x2
RSVPHop    - Ctype:1, Addr:192.168.36.1, LIH:3
TimeValue  - RefreshPeriod:30
SendTempl  - Sender:192.0.2.3, LspId:10
SendTSpec  - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
             MPU:20, MTU:1564
LabelReq   - IfType:General, L3ProtID:2048
RRO        - IpAddr:192.168.36.1, Flags:0x0
ERO        - IPv4Prefix 192.168.36.2/32, Strict
             IPv4Prefix 192.0.2.8/32, Loose
XRO        - IPv4Prefix: 192.0.2.4/32, Attribute: Node, LBit: Exclude
AdSpec     - General BreakBit:0, NumISHops:0, PathBwEstimate:0
                     MinPathLatency:4294967295, CompPathMTU:1564
             Controlled BreakBit:0
"
```

Node-ID Inclusion in the RESV Message

P-8 will be the MP for the bypass of ABR P-4 and PE-10 will be the MP for the bypass of ABR P-5. So P-8 and PE-10 need to include their node-ID in the RESV message, inside the Record Route Object (RRO).

```
*A:P-8# configure router rsvp node-id-in-rro include
*A:PE-10# configure router rsvp node-id-in-rro include
```

The default is **node-id-in-rro exclude**. As an example, the RESV message received on PLR node (PE-3) is shown below. The RRO contains the MP node (P-8) information in SLIL format:

```
*A:PE-3# debug router rsvp packet resv detail
18 2013/08/20 08:08:05.39 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: RESV Msg
Send RESV From:192.168.13.2, To:192.168.13.1
          TTL:255, Checksum:0x31f7, Flags:0x0
Session    - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
RSVPHop    - Ctype:1, Addr:192.168.13.2, LIH:3
TimeValue  - RefreshPeriod:30
Style      - SE
FlowSpec   - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
             MPU:20, MTU:1560, RSpecRate:0, RSpecSlack:0
FilterSpec - Sender:192.0.2.1, LspId:56852, Label:131071
  RRO      - …
             SystemIp:192.0.2.8, Flags:0x29
             Label:131070, Flags:0x1
             InterfaceIp:192.168.48.2, Flags:0x9
             Label:131070, Flags:0x1
             …
""
```

Bypass Configuration For ABR Protection

Since dynamic ABR protection is supported and used in this example, no explicit MBTs are configured to protect the ABRs. Each PLR first checks if an MBT tunnel exists between PLR and MP matching the constraints and protecting the ABR. If no MBT is available, the PLR will signal a bypass tunnel in a dynamic way towards MP node.

Figure 149 shows the two dynamic ABR node protections that are signalled for this LSP.

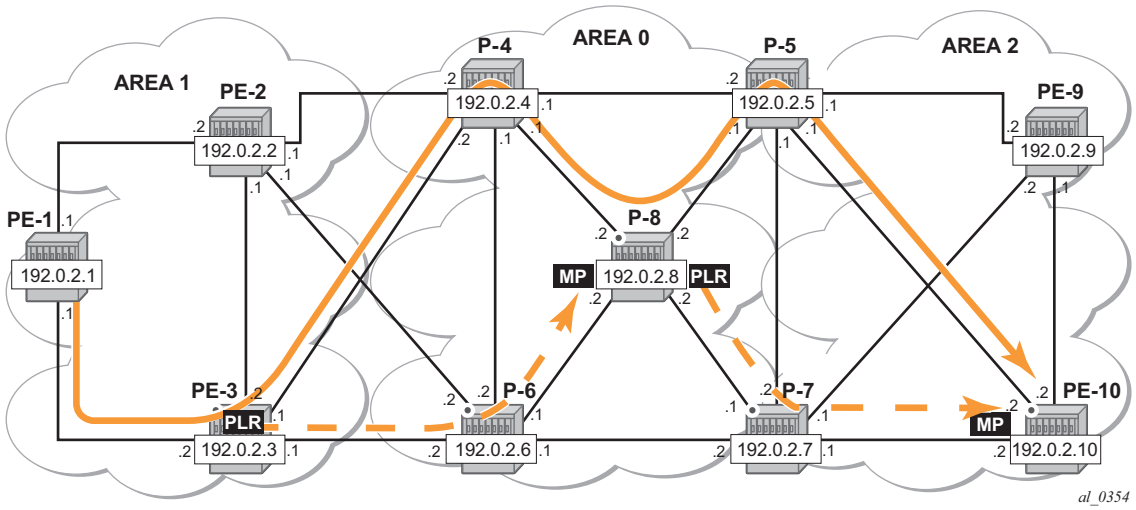

Figure 149: ABR Protection

Figure 150 shows the complete picture of all the FRR protections and indicates each node/link protection in the setup.

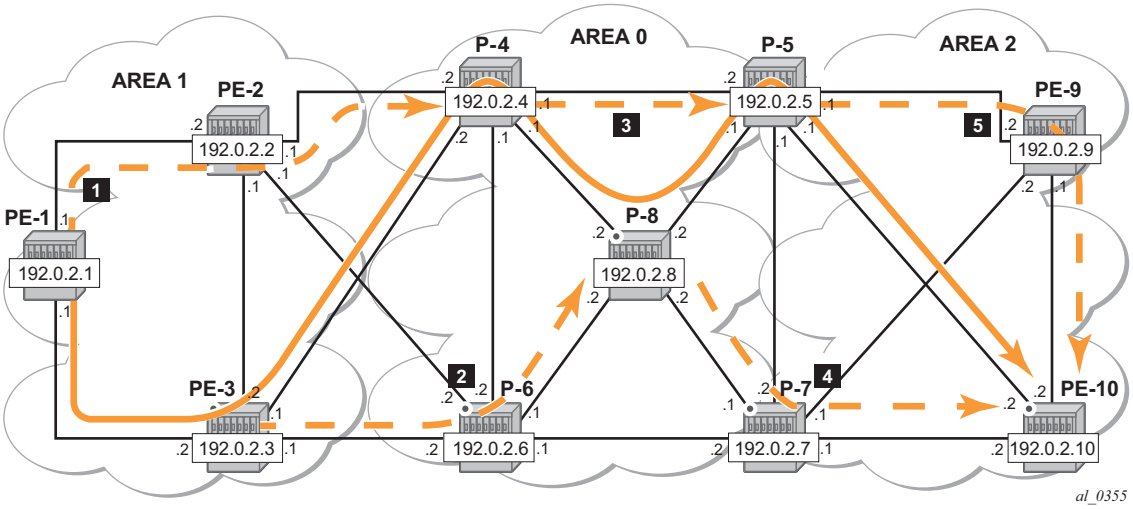

Figure 150: Protection of All Nodes/Links Along the LSP Path

This can be seen in the detailed show output of the LSP path:

```
*A:PE-1# show router mpls lsp "LSP-PE-1-PE-10" path detail
…
LSP Name     : LSP-PE-1-PE-10                    Path LSP ID : 56852
From         : 192.0.2.1                         To          : 192.0.2.10
Adm State    : Up                                Oper State  : Up
Path Name    : path-PE-10                        Path Type   : Primary
Path Admin   : Up                                Path Oper   : Up
…
Inter-area   : True
…
Actual Hops :
    192.168.13.1 (192.0.2.1) @ n       (1)       Record Label        : N/A
 -> 192.168.13.2 (192.0.2.3) @ n       (2)       Record Label        : 131071
 -> 192.168.34.2 (192.0.2.4) @ n       (3)       Record Label        : 131071
 -> 192.0.2.8 (192.0.2.8) @ n          (4)       Record Label        : 131070
 -> 192.168.48.2  @ n                  (4)       Record Label        : 131070
 -> 192.0.2.5 (192.0.2.5) @            (5)       Record Label        : 131071
 -> 192.168.58.1  @                    (5)       Record Label        : 131071
 -> 192.0.2.10 (192.0.2.10)                      Record Label        : 131071
 -> 192.168.105.2                                Record Label        : 131071
…
```

Note that there are two entries for P-8, P-5 and PE-10 in the ‘Actual Hops’ section in the previous output: one for the interface IP address and one for the system IP address. This is a consequence of configuring **node-id-in-rro include** on P-8, P-5 and PE-10.

Note: The **node-id-in-rro include** command is not mandatory for this example on ABR node P-5 but to be future save (for example, to cover cases where a new LSP is established in the network and P-5 acts as an MP node while the corresponding PLR node for that new LSP is in another area), this RSVP command can be enabled on all possible MP nodes in the network.

The details of the bypass tunnel can be shown with the following command:

```
*A:PE-3# show router mpls bypass-tunnel protected-lsp detail
===============================================================================
MPLS Bypass Tunnels (Detail)
===============================================================================
-------------------------------------------------------------------------------
bypass-node192.0.2.4
-------------------------------------------------------------------------------
To              : 192.0.2.8          State               : Up
Out I/F         : 1/1/4              Out Label           : 131070
Up Time         : 0d 01:41:04        Active Time         : n/a
Reserved BW     : 0 Kbps             Protected LSP Count : 1
Type            : Dynamic
Setup Priority  : 7                  Hold Priority       : 0
Class Type      : 0
Exclude Node    : 192.0.2.4          Inter-Area          : True
Computed Hops   :
    192.168.36.1(S)                  Egress Admin Groups : None
 -> 192.168.36.2(SA)                 Egress Admin Groups : None
 -> 192.0.2.8(L)                     Egress Admin Groups : None
Actual Hops     :
    192.168.36.1 (192.0.2.3)         Record Label        : N/A
 -> 192.168.36.2 (192.0.2.6)         Record Label        : 131070
 -> 192.0.2.8 (192.0.2.8)            Record Label        : 131069
 -> 192.168.68.2                     Record Label        : 131069

Protected LSPs -
LSP Name        : LSP-PE-1-PE-10::path-PE-10
From            : 192.0.2.1          To                  : 192.0.2.10
Avoid Node/Hop  : 192.0.2.4          Downstream Label    : 131070
Bandwidth       : 0 Kbps
```

The LSP could be further protected with one or more additional secondary paths, pre-signaled or not, but this is outside the scope of this example.

When a link or node failure occurs along the LSP path, FRR protection kicks in and end-to-end path re-optimization is executed: a PATHERR message is forwarded to the head-end. Upon receiving the PATHERR message the head-end calculates a new path.

Admin Groups

To support admin-groups for inter-area LSPs, the ingress node (PE-1) must propagate the admin-groups within the Session Attribute object (SA) of the PATH message so that the ABRs along the path receive the Admin Group restrictions they have to take into account when further expanding the ERO in the PATH message.

In Figure 151 the LSP path avoids the link between P-4 and P-8. This will be done by assigning admin group red to the link between P-4 and P-8 and then configuring the LSP to exclude the admin group red.

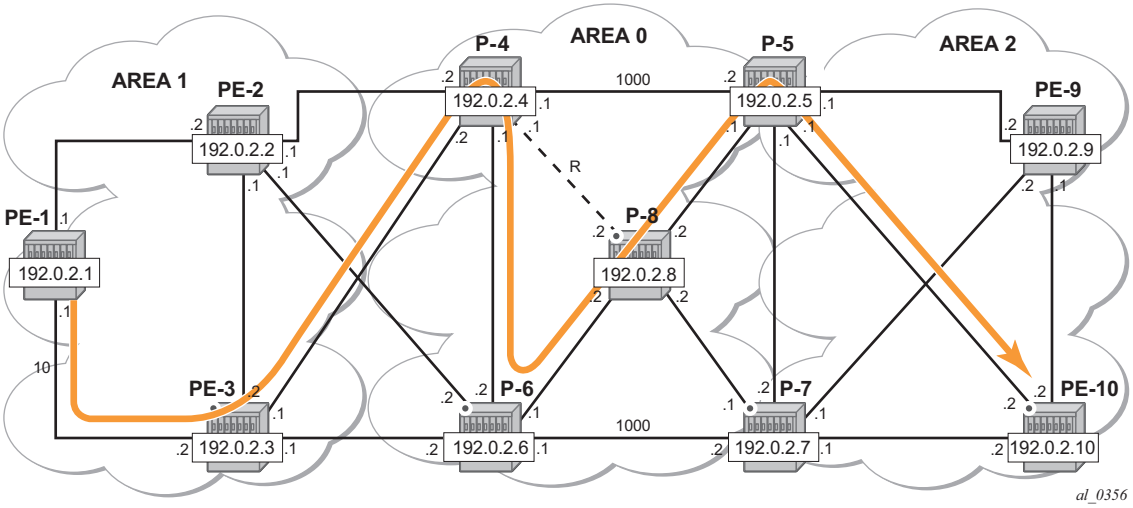

Figure 151: Admin Group Example

Admin Group Configuration

On P-4, configure admin group red and assign a group value to it (in the example value 11 is used, but this can be any value between 0 and 31). Assign admin group red to the link to P-8.

```
*A:P-4# configure router mpls
*A:P-4>config>router>mpls# admin-group "red" 11
*A:P-4>config>router>mpls# interface "int-P-4-P-8" admin-group "red"
```

Note that this admin-group configuration is required on all nodes in this example.

```
*A:Px# configure router mpls admin-group "red" 11
```

On PE-1, change the LSP configuration as follows:

```
*A:PE-1>config>router>mpls# info
----------------------------------------------
            admin-group "red" 11
…
            path "path-PE-10"
                no shutdown
            exit
            lsp "LSP-PE-1-PE-10"
                to 192.0.2.10
                cspf
                exclude "red"
                propagate-admin-group
                fast-reroute facility
                exit
                no shutdown
            exit
            no shutdown
----------------------------------------------
```

Note the **propagate-admin-group** command is required to include the admin group properties in the SA object of the PATH message. Admin-group value is mapped to a 32-bitmap. In this example, value 11 means that the 12th bit is set, which means in binary 100000000000 or hex 0x800.

```
*A:PE-1# debug router rsvp packet path detail
48 2013/08/20 18:41:05.89 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.1, To:192.0.2.10
          TTL:255, Checksum:0x80db, Flags:0x0
Session    - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
SessAttr   - Name:LSP-PE-1-PE-10::path-PE-10
             SetupPri:7, HoldPri:0, Flags:0x17
             Ctype:RA, ExcAny:0x800, IncAny:0x0, IncAll:0x0
RSVPHop    - Ctype:1, Addr:192.168.13.1, LIH:3
TimeValue  - RefreshPeriod:30
SendTempl  - Sender:192.0.2.1, LspId:56858
SendTSpec  - Ctype:QOS, CDR:0.000 bps, PBS:0.000 bps, PDR:infinity
```

```
              MPU:20, MTU:1564
LabelReq    - IfType:General, L3ProtID:2048
RRO         - IpAddr:192.168.13.1, Flags:0x0
ERO         - IPv4Prefix 192.168.13.2/32, Strict
              IPv4Prefix 192.168.34.2/32, Strict
              IPv4Prefix 192.0.2.10/32, Loose
FRRObj      - SetupPri:7, HoldPri:0, HopLimit:16, BW:0.000 bps, Flags:0x2
              ExcAny:0x0, IncAny:0x0, IncAll:0x0
"
```

The two sets of output below show that when P-4 expands the ERO it now excludes the link to node P-8 for the path calculation and the path is setup through P-6, P-8 and P-5.

```
*A:P-4# debug router rsvp packet path detail
9 2013/08/20 11:16:08.94 UTC MINOR: DEBUG #2001 Base RSVP
"RSVP: PATH Msg
Send PATH From:192.0.2.1, To:192.0.2.10
          TTL:253, Checksum:0xef94, Flags:0x0
Session     - EndPt:192.0.2.10, TunnId:2, ExtTunnId:192.0.2.1
SessAttr    - Name:LSP-PE-1-PE-10::path-PE-10
              SetupPri:7, HoldPri:0, Flags:0x17
              Ctype:RA, ExcAny:0x800, IncAny:0x0, IncAll:0x0
…
ERO         - IPv4Prefix 192.168.46.2/32, Strict
              IPv4Prefix 192.168.68.2/32, Strict
              IPv4Prefix 192.168.58.1/32, Strict
              IPv4Prefix 192.0.2.10/32, Loose
…
```

```
*A:PE-1# show router mpls lsp "LSP-PE-1-PE-10" path detail
…
LSP Name    : LSP-PE-1-PE-10                    Path LSP ID : 56858
From        : 192.0.2.1                         To          : 192.0.2.10
Adm State   : Up                                Oper State  : Up
Path Name   : path-PE-10                        Path Type   : Primary
Path Admin  : Up                                Path Oper   : Up
…
Actual Hops :
    192.168.13.1 (192.0.2.1) @ n                Record Label        : N/A
 -> 192.168.13.2 (192.0.2.3) @ n                Record Label        : 131071
 -> 192.168.34.2 (192.0.2.4) @ n                Record Label        : 131071
 -> 192.0.2.6 (192.0.2.6) @ n                   Record Label        : 131071
 -> 192.168.46.2  @ n                           Record Label        : 131071
 -> 192.0.2.8 (192.0.2.8) @                     Record Label        : 131071
 -> 192.168.68.2  @                             Record Label        : 131071
 -> 192.0.2.5 (192.0.2.5) @                     Record Label        : 131071
 -> 192.168.58.1  @                             Record Label        : 131071
 -> 192.0.2.10 (192.0.2.10)                     Record Label        : 131071
 -> 192.168.105.2                               Record Label        : 131071
…
```

Shared Risk Link Groups (SRLG)

Shared Risk Link Groups are also supported in the context of inter-area TE LSPs. SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail as well. Links in the group have a shared risk.

The MPLS TE SRLG feature enhances backup tunnel path selection so that a backup tunnel avoids using links that are in the same SRLG as interfaces the backup tunnel is protecting.

Consider the setup in Figure 152, where an inter-area LSP is setup from PE-1 to PE-10 and the path goes through P-8 because of a lower IGP metric. To protect against a node failure of P-8, P-4 (PLR) would normally setup an FRR backup directly to P-5 (MP), because of the lower IGP metric (P-4 to P-5:1000) compared to the IGP traffic via P-6 (P-4 to P-6 to P-7 to P-5:1200).

However, imagine that in this setup the P-4 <=> P-5 link and the P-4 <=> P-8 links are part of the same transmission bundle. In this case a cut of that fiber bundle will bring down both the primary and the backup path.

This can be avoided by configuring these two links in the same SRLG group and enabling srlg-frr strict on P-4. In that case the backup will be setup via P-6 as indicated by the dotted line in Figure 152

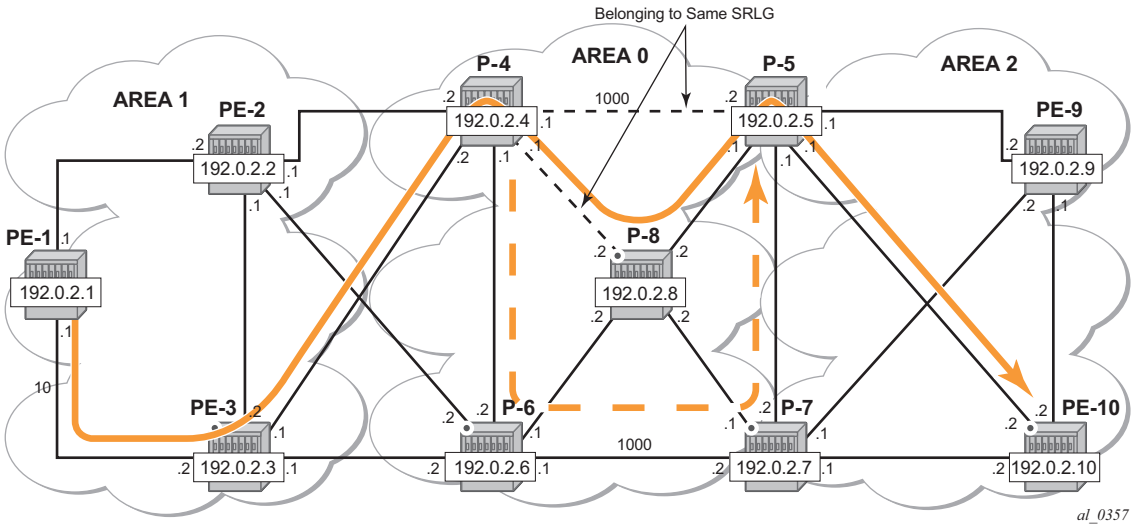

Figure 152: Share Risk Link Groups

SRLG Configuration

On P-4 configure an SRLG group and add the link to P-5 and the link to P-8 to this SRLG group and enable srlg-frr strict.

```
*A:P-4# configure router mpls
*A:P-4>config>router>mpls# srlg-group "bundle-red" value 1
*A:P-4>config>router>mpls# interface "int-P-4-P-5" srlg-group "bundle-red"
*A:P-4>config>router>mpls# interface "int-P-4-P-8" srlg-group "bundle-red"
*A:P-4>config>router>mpls#
*A:P-4>config>router>mpls# srlg-frr strict
```

Note that the srlg-group configuration is required on all nodes that use srlg groups and on the ABR used by the inter-area TE LSP.

```
*A:Px# configure router mpls  srlg-group "bundle-red" value 1
```

LSP Configuration

Remove the admin group restriction from the LSP.

```
*A:P-4# configure router mpls
*A:PE-1>config>router>mpls# info
----------------------------------------------
…
            lsp "PE-10"
                to 192.0.2.10
                cspf
                cspf-to-first-loose
                fast-reroute facility
                exit
                primary "path-PE-10"
                exit
                no shutdown
```

Now check the LSP path on PE-1 and verify that FRR protection is in place.

```
*A:PE-1# show router mpls lsp "LSP-PE-1-PE-10" path detail
===============================================================================
MPLS LSP LSP-PE-1-PE-10 Path  (Detail)
===============================================================================
Legend :
    @ - Detour Available              # - Detour In Use
    b - Bandwidth Protected           n - Node Protected
    s - Soft Preemption
    S - Strict                        L - Loose
    A - ABR
===============================================================================
-------------------------------------------------------------------------------
LSP LSP-PE-1-PE-10 Path path-PE-10
-------------------------------------------------------------------------------
```

```
LSP Name    : LSP-PE-1-PE-10                    Path LSP ID : 56864
From        : 192.0.2.1                         To          : 192.0.2.10
Adm State   : Up                                Oper State  : Up
Path Name   : path-PE-10                        Path Type   : Primary
Path Admin  : Up                                Path Oper   : Up
OutInterface: 1/1/1                             Out Label   : 131071
Path Up Time: 0d 00:15:26                       Path Dn Time: 0d 00:00:00
Retry Limit : 0                                 Retry Timer : 30 sec
RetryAttempt: 0                                 NextRetryIn : 0 sec

Adspec      : Disabled                          Oper Adspec : Disabled
CSPF        : Enabled                           Oper CSPF   : Enabled
Least Fill  : Disabled                          Oper LeastF*: Disabled
FRR         : Enabled                           Oper FRR    : Enabled
FRR NodePro*: Enabled                           Oper FRR NP : Enabled
FR Hop Limit: 16                                Oper FRHopL*: 16
FR Prop Adm*: Disabled                          Oper FRProp*: Disabled
Prop Adm Grp: Disabled                          Oper PropAG : Disabled
Inter-area  : True

Neg MTU     : 1560                              Oper MTU    : 1560
Bandwidth   : No Reservation                    Oper Bw     : 0 Mbps
Hop Limit   : 255                               Oper HopLim*: 255
Record Route: Record                            Oper RecRou*: Record
Record Label: Record                            Oper RecLab*: Record
SetupPriori*: 7                                 Oper SetupP*: 7
Hold Priori*: 0                                 Oper HoldPr*: 0
Class Type  : 0                                 Oper CT     : 0
Backup CT   : None
MainCT Retry: n/a
    Rem     :
MainCT Retry: 0
    Limit   :
Include Grps:                                     Oper InclGr*:
None                                            None
Exclude Grps:                                     Oper ExclGr*:
None                                            None

Adaptive    : Enabled                             Oper Metric : 110
Preference  : n/a
Path Trans  : 31                                  CSPF Queries: 17
Failure Code: noError                             Failure Node: n/a
ExplicitHops:
    No Hops Specified
Actual Hops :
    192.168.13.1 (192.0.2.1) @ n                  Record Label        : N/A
 -> 192.168.13.2 (192.0.2.3) @ n                  Record Label        : 131071
 -> 192.168.34.2 (192.0.2.4) @ n                  Record Label        : 131071
 -> 192.0.2.8 (192.0.2.8) @ n                     Record Label        : 131071
 -> 192.168.48.2  @ n                             Record Label        : 131071
 -> 192.0.2.5 (192.0.2.5) @                       Record Label        : 131071
 -> 192.168.58.1  @                               Record Label        : 131071
 -> 192.0.2.10 (192.0.2.10)                       Record Label        : 131070
 -> 192.168.105.2                                 Record Label        : 131070
ComputedHops:
    192.168.13.1(S)
 -> 192.168.13.2(S)
 -> 192.168.34.2(SA)
 -> 192.0.2.10(L)
```

```
ResigEligib*: False
LastResignal: n/a                                CSPF Metric : 110
===============================================================================
```

On P-4 check the SRLG configuration and verify that the backup is setup via P-6 rather than via P-5.

```
*A:P-4# show router mpls srlg-group
===============================================================================
MPLS Srlg Groups
===============================================================================
Group Name                      Group Value   Interfaces
-------------------------------------------------------------------------------
bundle-red                      1             int-P-4-P-5
                                              int-P-4-P-8
-------------------------------------------------------------------------------
No. of Groups: 1
===============================================================================


*A:P-4# show router mpls bypass-tunnel protected-lsp detail
===============================================================================
MPLS Bypass Tunnels (Detail)
===============================================================================
-------------------------------------------------------------------------------
bypass-node192.0.2.8
-------------------------------------------------------------------------------
To              : 192.168.57.1       State               : Up
Out I/F         : 1/1/5              Out Label           : 131068
Up Time         : 0d 00:21:37        Active Time         : n/a
Reserved BW     : 0 Kbps             Protected LSP Count : 1
Type            : Dynamic
Setup Priority  : 7                  Hold Priority       : 0
Class Type      : 0
Exclude Node    : None               Inter-Area          : False
Computed Hops   :
    192.168.46.1(S)                  Egress Admin Groups : None
 -> 192.168.46.2(S)                  Egress Admin Groups : None
 -> 192.168.67.2(S)                  Egress Admin Groups : None
 -> 192.168.57.1(S)                  Egress Admin Groups : None
Actual Hops     :
    192.168.46.1 (192.0.2.4)         Record Label        : N/A
 -> 192.168.46.2 (192.0.2.6)         Record Label        : 131068
 -> 192.168.67.2 (192.0.2.7)         Record Label        : 131069
 -> 192.168.57.1 (192.0.2.5)         Record Label        : 131070

Protected LSPs -
LSP Name        : LSP-PE-1-PE-10::path-PE-10
From            : 192.0.2.1          To                  : 192.0.2.10
Avoid Node/Hop  : 192.0.2.8          Downstream Label    : 131071
Bandwidth       : 0 Kbps
===============================================================================
```

Conclusion

Inter-area TE P2P LSPs can be setup based on ERO expansion. With this feature the head-end does a partial CSPF calculation to its local ABR. This ABR, on receiving a PATH message with a loose hop ERO, does a partial CSPF calculation to the next ABR or full CSPF to reach the final destination.

FRR protection within the area is available. FRR node protection of the ABR is possible through an MBT on the PLR (node just upstream of the ABR) to the MP (node just downstream of the ABR) or through a dynamically signaled bypass tunnel on the PLR. Dynamic ABR node protection requires that the node-ID of the MP node is propagated in the RESV message and that an XRO object is included in the bypass PATH message which makes it possible for the ABR to calculate a path to MP node.

TE features like BW, path prioritization, path pre-emption, graceful shutdown are supported, as well as propagation of the session attribute with affinity along the LSP path (admin groups) and SRLG.