

IGP Shortcuts

In This Chapter

This section provides information about IGP shortcuts.

Topics in this section include:

- [Applicability on page 982](#)
- [Overview on page 983](#)
- [Configuration on page 986](#)
- [Conclusion on page 1037](#)

Applicability

This example is applicable to the 7950 XRS, the 7750 and 7710 SR series and the 7450 platforms when the feature is not related to BGP and was tested on release 12.0.R3. There are no other prerequisites for this configuration.

Overview

Interior Gateway Protocols (IGPs) are routing protocols that operate inside an AS (Autonomous System). An AS is a network domain that is managed under a single administration. Because the scope of operation of an IGP is usually within an AS, IGPs are also called intra-AS protocols. The purpose of an IGP is to provide reachability information to destination nodes that are inside the domain. IGPs can be one or more of a variety of protocols, including routing protocols such as RIP version 1 or 2, OSPF, and IS-IS.

IGPs such as OSPF and IS-IS are link-state protocols that use an SPF algorithm to compute the shortest path tree to all nodes in a network. The results of such computations can be represented by the destination node, next-hop address, and output interface, where the output interface is a physical interface. Optionally, MPLS (Multiprotocol Label Switching) LSPs (Label Switched Paths) can be included in the SPF algorithm on the node performing the calculations, as LSPs behave as logical interfaces directly connected to remote nodes in the network. Because the SPF algorithm treats the LSPs in the same way as a physical interface (being a potential output interface), the computation results could be to select a destination node together with an output LSP, using the LSP as a shortcut through the network to the destination node.

Figure 140 shows a normal SPF tree sourced by PE-1 (Provider Edge-1).




Figure 140: Normal SPF Tree Sourced by PE-1

If there is an LSP that connects PE-1 to PE-5, and IGP shortcuts are configured on PE-1, the SPF tree will be as shown in [Figure 141](#).




Figure 141: SPF Tree Sourced by PE-1 Using LSP Shortcuts

IGP shortcuts are enabled on a per router basis; SPF computations are independent and irrelevant to other routers, so there is no need to enable shortcuts on every single router.

The network topology used in this example is displayed in [Figure 142](#). The setup consists of six 7750 service routers. There is a single AS and a single IGP area. The following configuration tasks should be completed first:

- IS-IS or OSPF on all interfaces within the AS (configuration has been done using IS-IS but using OSPF shows exactly the same behavior).
- LDP and RSVP on all interfaces within the AS.

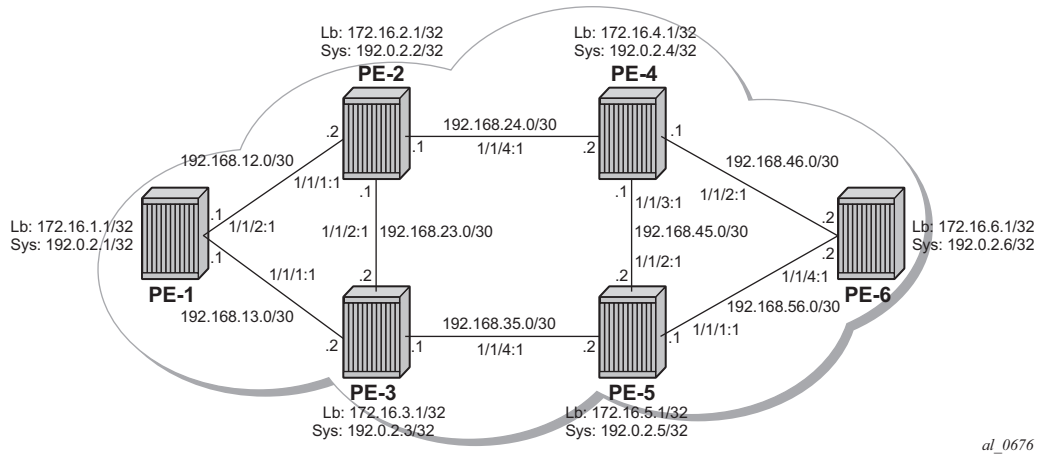


Figure 142: Tested Network Topology

Note: In all figures, **Lb** stands for Loopback and **Sys** stands for the system IP addresses.

Configuration

The first step is to configure the IGP (IS-IS) on all nodes, where IS-IS redistributes route reachability to all routers. To facilitate IS-IS configuration, all routers are L2-L1 capable within the same IS-IS area-id so there is only a single topology area in the network (all routers share the same topology). Traffic engineering is enabled on the IGP as it is a requirement for RSVP. The metric is using the default values: since no reference-bandwidth command is used, the default metric of 10 is applicable on all interfaces. The configuration for PE-2 is displayed below.

```
*A:PE-2#configure router
#-----
echo "IP Configuration"
#-----
    interface "int-PE-2-PE-1"
        address 192.168.12.2/30
        port 1/1/1:1
    exit
    interface "int-PE-2-PE-3"
        address 192.168.23.1/30
        port 1/1/2:1
    exit
    interface "int-PE-2-PE-4"
        address 192.168.24.1/30
        port 1/1/4:1
    exit
    interface "system"
        address 192.0.2.2/32
    exit
#-----
echo "ISIS Configuration"
#-----
    isis
        level-capability level-1/2  ## default
        area-id 49.0001
        traffic-engineering
        interface "system"
            passive
        exit
        interface "int-PE-2-PE-1"
            interface-type point-to-point
        exit
        interface "int-PE-2-PE-3"
            interface-type point-to-point
        exit
        interface "int-PE-2-PE-4"
            interface-type point-to-point
        exit
    exit
```

The configuration for the other nodes is very similar. The IP addresses can be derived from [Figure 142](#).

The GRT (Global Route Table) for PE-2 is displayed below.

```
*A:PE-2# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.1/32                               Remote  ISIS    16h07m38s    15
    192.168.12.1                             10
192.0.2.2/32                               Local   Local   16h15m08s    0
    system                                     0
192.0.2.3/32                               Remote  ISIS    16h11m32s    15
    192.168.23.2                             10
192.0.2.4/32                               Remote  ISIS    16h03m59s    15
    192.168.24.2                             10
192.0.2.5/32                               Remote  ISIS    16h00m18s    15
    192.168.23.2                             20
192.0.2.6/32                               Remote  ISIS    16h03m59s    15
    192.168.24.2                             20
192.168.12.0/30                             Local   Local   16h14m56s    0
    int-PE-2-PE-1                             0
192.168.13.0/30                             Remote  ISIS    16h07m38s    15
    192.168.12.1                             20
192.168.23.0/30                             Local   Local   16h14m56s    0
    int-PE-2-PE-3                             0
192.168.24.0/30                             Local   Local   16h04m23s    0
    int-PE-2-PE-4                             0
192.168.35.0/30                             Remote  ISIS    16h00m18s    15
    192.168.23.2                             20
192.168.45.0/30                             Remote  ISIS    00h21m29s    15
    192.168.24.2                             20
192.168.46.0/30                             Remote  ISIS    16h03m59s    15
    192.168.24.2                             20
192.168.56.0/30                             Remote  ISIS    16h00m18s    15
    192.168.23.2                             30
-----
No. of Routes: 14
=====
```

LDP and RSVP Shortcuts

Interface Label Distribution Protocol (iLDP) is enabled on all interfaces (except system interfaces, which is not allowed) in all routers. The configuration on all nodes is similar and the IP addresses are derived from [Figure 142](#). Below is the configuration of PE-4.

```
A:PE-4>configure router ldp
  interface-parameters
    interface "int-PE-4-PE-2"
    exit
    interface "int-PE-4-PE-5"
    exit
    interface "int-PE-4-PE-6"
    exit
  exit
  targeted-session
  exit
```

With iLDP enabled, PE-4 establishes iLDP sessions with its directly connected neighbors, as shown below.

```
*A:PE-4# show router ldp session
=====
LDP Sessions
=====
Peer LDP Id          Adj Type   State           Msg Sent  Msg Recv  Up Time
-----
192.0.2.2:0         Link      Established    276       275       0d 00:12:23
192.0.2.5:0         Link      Established    2629      2632      0d 02:01:20
192.0.2.6:0         Link      Established    225       226       0d 00:10:06
-----
No. of Sessions: 3
=====
```

The following tunnel table shows that there is a Label Switched Path (LSP) to every other router. The reason is that the LDP label distribution mode is DU (Downstream Unsolicited) by default, originating label bindings for system addresses only (which are used by iLDP as transport address by default). The command also shows LSPs' preference (where the preference is 9 for LDP) and metric (metric is inherited from the IGP, each hop counts as a metric of 10), as shown below. The metric to destinations PE-1 and PE-3 is 20 because there are two hops in between (PE-4 is two hops away from PE-1 and PE-3).

```
*A:PE-4# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId  Pref  Nexthop      Metric
-----
192.0.2.1/32         ldp   MPLS    -        9     192.168.24.1  20
192.0.2.2/32         ldp   MPLS    -        9     192.168.24.1  10
192.0.2.3/32         ldp   MPLS    -        9     192.168.24.1  20
```



```

192.0.2.5/32      ldp  MPLS  -      9      192.168.45.2  10
192.0.2.6/32      ldp  MPLS  -      9      192.168.46.2  10
=====

```

In order to configure RSVP shortcuts, RSVP must be enabled on all interfaces where traffic engineering is required. For simplicity, RSVP is configured on all interfaces of the network, including system interfaces. The configuration for PE-6 is displayed below.

```

*A:PE-6#configure router
  mpls
    interface "system"
    exit
    interface "int-PE-6-PE-5"
    exit
    interface "int-PE-6-PE-4"
    exit
  exit
  rsvp
    interface "system"
    exit
    interface "int-PE-6-PE-5"
    exit
    interface "int-PE-6-PE-4"
    exit
    no shutdown
  exit

```

The configuration of the other nodes is similar. The IP addresses can be derived from [Figure 142](#). Because there are no RSVP LSPs configured yet, the tunnel-table has no RSVP LSPs and only contains LDP LSPs.

LDP Static Route (IP Tunneled in LDP Tunnel)

Using LDP LSP shortcuts for static route resolution enables forwarding of IPv4 packets over LDP LSPs instead of using a regular IP next-hop. In other words, the traffic to the resolved static routes is forwarded making use of an MPLS LDP LSP rather than plain IP.

The configuration defines a static route pointing to the destination PE (remote loopback, which is an indirect next hop in the example), and explicitly indicates that it should use LDP rather than IGP. Taking PE-1 and PE-6 as an example, two loopback interfaces are configured (172.16.X.1/32), where X = PE number, and a static-route is defined according to the explanation above. The following shows the configuration on PE-1.

```
*A:PE-1>config router
#-----
echo "IP Configuration"
#-----
    interface "int-PE-1-PE-2"
        address 192.168.12.1/30
        port 1/1/2:1
    exit
    interface "int-PE-1-PE-3"
        address 192.168.13.1/30
        port 1/1/1:1
    exit
    interface "loopback"
        address 172.16.1.1/32
        loopback
    exit
    interface "system"
        address 192.0.2.1/32
    exit
#-----
echo "Static Route Configuration"
#-----
    static-route 172.16.6.1/32 indirect 192.0.2.6 ldp disallow-igp
```

Looking at the GRT or FIB (Forwarding Database), there are two new entries corresponding to the two configured loopbacks. One entry is associated with protocol local (local loopback on the PE), and the other entry is protocol static, where the next hop is reached using a LDP LSP.

```
*A:PE-1# show router fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
172.16.1.1/32                             LOCAL
    172.16.1.1 (loopback)
172.16.6.1/32                             STATIC
    192.0.2.6 (Transport:LDP)
* Truncated info
```

The next output shows that a ping sourced by PE-1's loopback interface is able to reach PE-6's loopback, and traceroute demonstrates that the traffic is following the LDP LSP. The ping and traceroute commands cannot follow the IGP because the static-route command states that the IGP is disallowed when no LDP LSP towards PE-6 is available (also, the loopback interfaces are not enabled on IS-IS).

```
*A:PE-1# ping 172.16.6.1 source 172.16.1.1
PING 172.16.6.1 56 data bytes
64 bytes from 172.16.6.1: icmp_seq=1 ttl=64 time=4.08ms.
64 bytes from 172.16.6.1: icmp_seq=2 ttl=64 time=3.80ms.
64 bytes from 172.16.6.1: icmp_seq=3 ttl=64 time=4.77ms.
64 bytes from 172.16.6.1: icmp_seq=4 ttl=64 time=3.85ms.
64 bytes from 172.16.6.1: icmp_seq=5 ttl=64 time=3.83ms.

---- 172.16.6.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.80ms, avg = 4.07ms, max = 4.77ms, stddev = 0.371ms
*A:PE-1# traceroute 172.16.6.1 source 172.16.1.1
traceroute to 172.16.6.1 from 172.16.1.1, 30 hops max, 40 byte packets
 1  0.0.0.0  * * *
 2  0.0.0.0  * * *
 3  172.16.6.1 (172.16.6.1)    4.64 ms  9.87 ms  5.84 ms
```

With the traceroute command, there are three hops from PE-1 to PE-6. There is no information regarding IP for the first two hops because the traffic is encapsulated in an MPLS LDP. The reason why the hops are displayed even when there is an MPLS LSP tunnel is because by default, the SR router propagates (copies) the TTL (Time to Live) from the IP header in the MPLS header. This is known as uniform mode.

However, a service provider might not want to show how many MPLS hops (nodes) there are in their network if a **traceroute** command is executed from outside their network. To prevent internal hops being shown, no **propagate** commands are needed in the LDP configuration, as shown below. This is known as pipe mode.

```
*A:PE-1#config router ldp
      no shortcut-local-ttl-propagate
      no shortcut-transit-ttl-propagate
interface-parameters
      interface "int-PE-1-PE-2"
      exit
      interface "int-PE-1-PE-3"
      exit
exit
targeted-session
exit
```

Once TTL propagation is disabled, the hops are not displayed any longer when running the traceroute command.

```
*A:PE-1# traceroute 172.16.6.1 source 172.16.1.1
traceroute to 172.16.6.1 from 172.16.1.1, 30 hops max, 40 byte packets
 1  172.16.6.1 (172.16.6.1)    5.08 ms  4.73 ms  4.77 ms
```

RSVP Static Route (IP Tunneled in RSVP Tunnel)

Using RSVP LSP shortcuts for static route resolution enables forwarding of IPv4 packets over RSVP LSPs instead of using a regular IP next-hop. In other words, the traffic to the resolved static routes is forwarded making use of an MPLS RSVP LSP rather than plain IP.

The configuration defines a static route pointing to a destination PE (remote loopback, which is an indirect next hop in the example), and explicitly indicates that it should use RSVP rather than IGP. Taking PE-6 and PE-1 as an example, two loopback interfaces are configured (172.16.X.1/32), where X = PE number, and a static-route is defined according to the explanation above. The following shows the configuration on PE-6.

```
*A:PE-6>config router
#-----
echo "IP Configuration"
#-----
        interface "int-PE-6-PE-4"
            address 192.168.46.2/30
            port 1/1/2:1
        exit
        interface "int-PE-6-PE-5"
            address 192.168.56.2/30
            port 1/1/4:1
        exit
        interface "loopback"
            address 172.16.6.1/32
            loopback
        exit
        interface "system"
            address 192.0.2.6/32
        exit
#-----
echo "Static Route Configuration"
#-----
        static-route 172.16.1.1/32 indirect 192.0.2.1 rsvp-te disallow-igp
```

Also, an RSVP LSP needs to be configured with PE-1's system interface as the destination:

```
*A:PE-6>config router mpls
        interface "system"
            exit
            interface "int-PE-6-PE-5"
                exit
            interface "int-PE-6-PE-4"
                exit
            path "p"
                no shutdown
            exit
            lsp "LSP-PE-6-PE-1"
                to 192.0.2.1
                primary "p"
            exit
            no shutdown
        exit
```

```
no shutdown
```

Reviewing the LSP tunnel table, observe that there is an RSVP LSP created:

```
*A:PE-6>config>router# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId  Pref    Nexthop      Metric
-----
192.0.2.1/32         rsvp  MPLS   1         7      192.168.56.1  30
=====
```

Note that the default RSVP preference is 7 (preferred over that of LDP, which is 9) and the metric reflects that this LSP spans 3 hops (for a dynamic LSP not using CSPF, the metric is inherited from IGP). See [RSVP Shortcut for IGP Route Resolution on page 999](#) for more details about the metric applied in LSPs.

The RSVP LSP is used to resolve the indirect next hop (PE-1 system address) in the static route (the LSP used is identified with the Tunnel ID, in this case 1), hence the GRT is modified with the following entry:

```
*A:PE-6# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix          Type    Proto   Age      Pref
  Next Hop[Interface Name]      Metric
-----
172.16.1.1/32       Remote  Static  00h09m55s  5
    192.0.2.1 (tunneled:RSVP:1)  1
*Truncated info
-----
No. of Routes: 16
=====
```

As in the LDP shortcut with static route example, between PE-6 and PE-1, TTL propagation is disabled.

```
*A:PE-6#config router mpls
    no shortcut-local-ttl-propagate
    no shortcut-transit-ttl-propagate
```

RSVP Static Route (IP Tunneled in RSVP Tunnel)

The output is the following when running a traceroute:

```
*A:PE-6# traceroute 172.16.1.1 source 172.16.6.1
traceroute to 172.16.1.1 from 172.16.6.1, 30 hops max, 40 byte packets
 1 172.16.1.1 (172.16.1.1)    4.71 ms  24.7 ms  17.2 ms
```

Note that the two static routes that have been defined to use the LDP and RSVP shortcuts follow the static routes default values and have a preference of 5 and a metric of 1.

LDP Shortcut for IGP Route Resolution

Using LDP shortcuts for IGP route resolution enables forwarding of packets to IGP learned routes over an LDP LSP. The default is to disable the LDP shortcut across all interfaces in the node.

When LDP shortcuts are enabled, LDP populates the RTM (Route Table Manager) with next-hop entries corresponding to all prefixes for which it activated an LDP Forwarding Equivalence Class (FEC). For a given prefix, two route entries are populated in RTM. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always takes preference over the regular IP next-hop for forwarding user packets and specific control packets over a given outgoing interface to the route next-hop.

Once LDP has activated a FEC for a given prefix and programmed RTM, it also programs the ingress Tunnel Table in line card with the LDP tunnel information.

When an IPv4 packet is received on an ingress network interface, a subscriber IES (Internet Enhanced Service) interface, or a regular IES interface, the lookup of the packet by the ingress line card results in the packet being sent labeled with the label stack corresponding to the NHLFE (Next Hop Label Forwarding Entry) of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut. If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabeled. The activation of the FEC by LDP is done by performing an exact match with an IGP route prefix in RTM but it can also be done by performing a longest prefix-match with an IGP route in RTM if the aggregate-prefix-match option is enabled globally in LDP.

Handling of Control Packets

All control plane packets will not see the LDP shortcut route entry in RTM with the exception of the following control packets which will be forwarded over an LDP shortcut when enabled:

- A locally generated or in transit ICMP ping and UDP traceroute of an IGP route. The transit message appears as a user packet to the ingress LER node.
- A locally generated response to a received ICMP ping or UDP traceroute message.

All other control plane packets that require an RTM lookup and have knowledge of which destination is reachable over the LDP shortcut will continue to be forwarded over the IP next-hop route in RTM.

Handling of Multicast Packets

LDP shortcuts apply to unicast FEC types and are used for forwarding IP unicast packets in the data path. IP multicast packets forwarded over an mLDP P2MP LSP make use of a multicast FEC and thus cannot make use of the LDP unicast shortcut.

ECMP Considerations

When ECMP is enabled and multiple equal-cost next-hops exist for the IGP route, the ingress line card will spray the packets for this route based on the hashing routine supported for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying is performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP (Targeted LDP) neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both. This is as per ECMP for LDP in the existing implementation. When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix. Spraying across regular IP next-hops and LDP-shortcut next-hops concurrently is not supported.

Configuring IGP LDP shortcuts is straightforward, and only applies to the node where there is interest to provision the LDP shortcut. In this example, only PE-1 is provisioned with LDP shortcuts, as shown below.

```
*A:PE-1#config router ldp-shortcut
```

Now, all tunnel LSPs that resolve an IGP next hop will replace the IP next hops, as depicted in the following output:

```
*A:PE-1>config>router# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.1/32                               Local  Local   23h57m47s    0
  system                                   0
192.0.2.2/32                               Remote LDP    00h00m04s    9
  192.168.12.2 (tunneled)                  10
192.0.2.3/32                               Remote LDP    00h00m04s    9
  192.168.13.2 (tunneled)                  10
192.0.2.4/32                               Remote LDP    00h00m04s    9
  192.168.12.2 (tunneled)                  20
192.0.2.5/32                               Remote LDP    00h00m04s    9
  192.168.13.2 (tunneled)                  20
192.0.2.6/32                               Remote LDP    00h00m04s    9
  192.168.12.2 (tunneled)                  30
192.168.12.0/30                             Local  Local   23h57m33s    0
  int-PE-1-PE-2                             0
192.168.13.0/30                             Local  Local   02h40m29s    0
```



```

int-PE-1-PE-3
192.168.23.0/30 Remote ISIS 02h30m22s 15
    192.168.12.2 20
192.168.24.0/30 Remote ISIS 23h53m58s 15
    192.168.12.2 20
192.168.35.0/30 Remote ISIS 02h40m28s 15
    192.168.13.2 20
192.168.45.0/30 Remote ISIS 02h30m22s 15
    192.168.12.2 30
192.168.46.0/30 Remote ISIS 23h53m53s 15
    192.168.12.2 30
192.168.56.0/30 Remote ISIS 02h40m28s 15
    192.168.13.2 30

```

No. of Routes: 14

*A:PE-1>config>router# show router fib 1

=====

FIB Display

=====

Prefix	NextHop	Protocol
192.0.2.1/32		LOCAL
192.0.2.1 (system)		
192.0.2.2/32		LDP
192.0.2.2 (Transport:LDP)		
192.0.2.3/32		LDP
192.0.2.3 (Transport:LDP)		
192.0.2.4/32		LDP
192.0.2.4 (Transport:LDP)		
192.0.2.5/32		LDP
192.0.2.5 (Transport:LDP)		
192.0.2.6/32		LDP
192.0.2.6 (Transport:LDP)		
192.168.12.0/30		LOCAL
192.168.12.0 (int-PE-1-PE-2)		
192.168.13.0/30		LOCAL
192.168.13.0 (int-PE-1-PE-3)		
192.168.23.0/30		ISIS
192.168.12.2 (int-PE-1-PE-2)		
192.168.24.0/30		ISIS
192.168.12.2 (int-PE-1-PE-2)		
192.168.35.0/30		ISIS
192.168.13.2 (int-PE-1-PE-3)		
192.168.45.0/30		ISIS
192.168.12.2 (int-PE-1-PE-2)		
192.168.46.0/30		ISIS
192.168.12.2 (int-PE-1-PE-2)		
192.168.56.0/30		ISIS
192.168.13.2 (int-PE-1-PE-3)		

Total Entries : 14

LDP Shortcut for IGP Route Resolution

Applying LDP IGP shortcuts only on PE-1 implies that IP traffic from PE-1 to any of the system addresses of the rest of nodes will use the LDP shortcut, however, the traffic replied from any PE back to PE-1 will be native IP since IGP shortcuts have not been provisioned in the other nodes.

RSVP Shortcut for IGP Route Resolution

Using RSVP LSP shortcuts when resolving IGP routes enables forwarding of packets to IGP learned routes over an RSVP LSP. The use of RSVP shortcut for resolving IGP routes is enabled at the IS-IS (or OSPF) routing protocol level or at the LSP level, and instructs IS-IS and OSPF to include RSVP LSPs originating on this node and terminating on the system address (router-id) of a remote node and considers them as direct links. RSVP LSPs with a destination address corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF.

By default, **rsvp-shortcut** is disabled in all IGP instances.

- RSVP LSPs are included in the IGP SPF computation with the following characteristics:
- RSVP LSP is modeled as a point-to-point link IP interface and its metric is used in the computation of the shortest path of IGP routes

Next-hop and interface do include the NHLFE of the shortcut LSP when the IGP path cost using the RSVP LSP is the best.

Shortcuts are not used when the destination RSVP LSP is in a different IGP area. In addition, IGP adjacencies across an RSVP LSP are not supported.

The next output shows the configuration commands:

```
*A:PE>config>router>isis/ospf
- rsvp-shortcut
```

Note that the configuration can be done at the IGP level or per LSP level. When `rsvp-shortcut` is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the command

```
*A:PE>config>router>mpls>lsp#no igp-shortcut
```

As RSVP shortcuts can coexist with LDP shortcuts or IP next hops, SPF computation and path selection follows the procedures in RFC 3906:

- SPF picks the RSVP shortcut next-hop if there is an RSVP LSP directly to that address regardless of the path cost compared to the IGP next-hop.
- SPF picks the RSVP shortcut next-hop or the IGP next-hop based on path lowest cost if there is an IGP path to the prefix that does not go via the tail-end of the LSP.
- If the IGP next-hop is picked, then it can be an LDP shortcut next-hop or a regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop.

Handling of Control Packets

All control plane packets requiring an RTM lookup and whose destination is reachable over the RSVP shortcut are forwarded over the shortcut. This is because RTM keeps a single route entry for each prefix except if there is ECMP over different outgoing interfaces. Interface bound control packets are not impacted by the RSVP shortcut since RSVP LSPs with a destination address different than the router-id are not included by IGP in its SPF calculation.

Important note: RSVP shortcuts for IGP shortcut resolution should only be used with CSPF LSPs and/or with fully explicit path non CSPF LSP. RSVP hop-by-hop Path messages will try to use the shortcut and consequently LSPs without CSPF enabled, or that use a loose/empty hop path, will not come up. However, LSPs with CSPF enabled or using a strict hop path will come up. This is because in the former case the RTM lookup to get the next hop results in using the shortcut and so the path messages are sent directly to the destination of the LSP, where they are dropped. With CSPF enabled, the next-hop (and the entire path) is provided by CSPF and the path messages are sent unlabeled to the directly connected neighbor which corresponds to the next-hop of the destination of the LSP. Similar processing occurs if a strict hop path is used, as is the case in the example below.

Handling of Multicast Packets

IP multicast packets cannot be forwarded over an RSVP shortcut, they can only be forwarded over an RSVP P2MP LSP. However, as RSVP shortcut routes appear in RTM and are seen by all applications when they are the best route. When the reverse path forwarding (RPF) check for the source of the multicast packet matches an RSVP shortcut route, the check will pass if both the RSVP shortcut and the multicast-import options are enabled in the IGP, as shown below, as the RTM is populated with next hops only and not with tunnels (RPFs will fail for source prefixes resolved to a tunnel NH).

```
A:PE-2>config>router>isis# multicast-import
- multicast-import [both]
- multicast-import [ipv4]
- multicast-import [ipv6]
- no multicast-import [both]
- no multicast-import [ipv4]
- no multicast-import [ipv6]

<ipv4>          : keyword
<ipv6>          : keyword
<both>         : keyword
```

The unicast RTM can still make use of the tunnel next-hop for the same prefix. SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

ECMP Considerations

When ECMP is enabled and multiple equal-cost paths exist for the route over a set of tunnel next-hops based on the hashing routine supported for IPv4 packets, there are two possibilities:

- Destination is tunnel-endpoint: the system selects the tunnel with lowest tunnel ID (IP next-hop is never used).
- Destination is different from the tunnel endpoint: it selects tunnel endpoints when the LSP metric is not greater than the IGP cost and it prefers tunnel endpoint over IP next-hop.

Note that ECMP is not performed across both the IP and tunnel next-hops.

RSVP Shortcuts Configuration

Configuring RSVP LSP shortcuts is straightforward, and only applies to the node where there is interest to provision the RSVP shortcut. Two LSPs, from PE-6 to PE-1 and from PE-1 to PE-6, with strict hops, are provisioned according to [Figure 143](#), as shown below.




Figure 143: LSPs Between PE-1 and PE-6

The configuration on PE-1 and PE-6 is similar (replacing the IP addresses), so only the configuration for PE-6 is shown:

```
*A:PE-6#config router isis
  area-id 49.0001
  traffic-engineering
  rsvp-shortcut
  interface "system"
    passive
  exit
  interface "int-PE-6-PE-5"
    interface-type point-to-point
  exit
  interface "int-PE-6-PE-4"
    interface-type point-to-point
  exit
exit

*A:PE-6#config router mpls
  path "PE-1"
    hop 1 192.0.2.5 strict
    hop 2 192.0.2.3 strict
    hop 3 192.0.2.2 strict
    hop 4 192.0.2.1 strict
    no shutdown
  exit
```

```

lsp "LSP-PE-6-PE-1"
  to 192.0.2.1
  primary "PE-1"
  exit
  no shutdown
exit
no shutdown
exit

```

The GRT output shows the change in the next hop, using an RSVP shortcut:

```

*A:PE-6# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type  Proto  Age      Pref
  Next Hop[Interface Name]                Metric
-----
192.16.6.1/32                             Local  Local  21h40m00s  0
  loopback
192.0.2.1/32                             Remote ISIS   00h00m07s 15
  192.0.2.1 (tunneled:RSVP:1)             16777215
192.0.2.2/32                             Remote  ISIS   00h00m07s  15
  192.168.46.1
192.0.2.3/32                             Remote  ISIS   01d22h34m  15
  192.168.56.1
192.0.2.4/32                             Remote  ISIS   01d22h42m  15
  192.168.46.1
192.0.2.5/32                             Remote  ISIS   07d22h10m  15
  192.168.56.1
192.0.2.6/32                             Local  Local  07d22h12m  0
  system
192.168.12.0/30                           Remote  ISIS   00h00m07s  15
  192.168.46.1
192.168.13.0/30                           Remote  ISIS   00h00m07s  15
  192.168.56.1
192.168.23.0/30                           Remote  ISIS   00h00m07s  15
  192.168.46.1
192.168.24.0/30                           Remote  ISIS   01d22h42m  15
  192.168.46.1
192.168.35.0/30                           Remote  ISIS   07d22h10m  15
  192.168.56.1
192.168.45.0/30                           Remote  ISIS   01d01h14m  15
  192.168.46.1
192.168.46.0/30                           Local  Local  03d08h47m  0
  int-PE-6-PE-4
192.168.56.0/30                           Local  Local  07d22h11m  0
  int-PE-6-PE-5
-----
No. of Routes: 15

```

The RSVP LSP in the output has a metric of 16777215, the LSP administrative metric matches the maximum value allowed for an IS-IS link using the wide-metric (24-bit value with a range of [0—16777215]). The following metric rules apply:

RSVP Shortcut for IGP Route Resolution

- A dynamic strict path non-CSPF LSP has the maximum metric (16777215).
- A dynamic CSPF LSP has a metric equal to the cumulative IGP cost.
 - ☞ If the user enabled the use of the TE metric on this LSP (configure router mpls lsp cspf use-te-metric), then the metric for the LSP is the maximum (16777215).
 - ☞ If the user enabled the use of the TE metric on this LSP (configure router mpls lsp cspf use-te-metric) and provisioned a specific metric on the lsp (configure router mpls lsp metric <metric>:<0..16777215>), then the metric for the LSP is the one provisioned. Note that when configuring the metric of an LSP, the parameter “use-te-metric” is not required.
- A static LSP has a maximum metric (16777215).
- Manual and dynamic bypass LSPs have the maximum metric (16777215).

Note: The RSVP shortcuts section detailed the importance of the LSP metric when using CSPF LSPs or when importing RSVP tunnel links into the IGP. The LSP metric can be inherited from the IGP, or can be manually modified by configuring a specific LSP metric or relative-metric offset. As IP and LDP FECs resolve to RSVP LSPs when the metric is equal or lower compared to the regular routing metric, configuring a specific static LSP metric (lower than the IGP metric) or relative-metric offset is strongly recommended when using RSVP shortcuts, so that the GRT and LDP FEC resolution will always prefer to use RSVP LSP shortcuts when the CSPF path computation is not using the shortest path.

For the example above, first rule applies.

Advertising RSVP LSP Tunnel Links in the IGP: Forwarding Adjacency Feature

If configured, an RSVP LSP can also be advertised into the IGP similar to regular links so that other routers in the network can include it into their SPF computations. The forwarding adjacency feature can be enabled independently from the RSVP shortcut feature in CLI. If both are configured for a given IGP instance, the forwarding adjacency takes precedence. An RSVP LSP must exist in the reverse direction in order for the advertised link to pass the bi-directional link check and be usable by other routers in the network. However, this is not required for the node which originates the LSP. The LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs as per RFC 3906.

Reusing the RSVP IGP shortcuts set up previously (PE-1 and PE-6 RSVP IGP shortcut example according to [Figure 143](#)), the outcome is a route linked with an RSVP LSP as next hop, as seen below:

```
*A:PE-6# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.1/32                               Remote ISIS   00h00m07s   15
  192.0.2.1 (tunneled:RSVP:1)              16777215
* Truncated info                            0
-----
No. of Routes: 15
```

The route tunneled through RSVP has a metric of 16777215, so it is not used by PE-6 GRT to reach any other routes since the metric is very high. After enabling the forwarding adjacency feature (tunnel links) to use shortcuts in the configuration, PE-1 and PE-6 have a direct connection through the RSVP LSP (as a virtual link). This configuration command must be executed in both routers, although for simplicity only PE-6 is displayed:

```
*A:PE-6#config router isis
  advertise-tunnel-link
```

Once the shortcut is advertised by IS-IS, the route will disappear from the RTM as the metric of the shortcut is greater than the IGP cost.

```
A:PE-6>config>router# show router route-table 192.0.2.1/32
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                        Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.1/32                               Remote ISIS   00h00m01s   15
  192.168.46.1                              30
```

RSVP Shortcut for IGP Route Resolution

```
-----  
No. of Routes: 1
```

If the LSP is reconfigured to use a metric equal to or smaller than the IGP cost, the router PE-6 will use the RSVP shortcut again. In the example, the LSP is reconfigured with a metric of 30:

```
*A:PE-6#config router mpls  
  lsp "LSP-PE-6-PE-1" metric 30
```

Now the shortcut shows up as the preferred next hop to reach PE-1 from PE-6.

```
A:PE-6>config>router# show router route-table 192.0.2.1  
-----  
Route Table (Router: Base)  
-----  
Dest Prefix[Flags]                               Type   Proto   Age           Pref  
  Next Hop[Interface Name]                       Metric  
-----  
192.0.2.1/32                                     Remote ISIS   00h00m02s  15  
  192.0.2.1 (tunneled:RSVP:1)                   30  
-----  
No. of Routes: 1
```

As explained earlier, this could be combined together with ECMP, so if ECMP is configured to 2, the system shows the two equal cost paths.

```
A:PE-6>config>router# ecmp 2  
A:PE-6>config>router# show router route-table 192.0.2.1  
-----  
Route Table (Router: Base)  
-----  
Dest Prefix[Flags]                               Type   Proto   Age           Pref  
  Next Hop[Interface Name]                       Metric  
-----  
192.0.2.1/32                                     Remote ISIS   00h00m00s  15  
  192.0.2.1 (tunneled:RSVP:1)                   30  
192.0.2.1/32                                     Remote ISIS   00h00m00s  15  
  192.168.46.1                                   30  
-----  
No. of Routes: 2
```

Checking GRT on PE-4, it displays the route to reach PE-1 (192.0.2.1/32) with a metric of 20 via PE-2 as next-hop. Although now PE-6 is announcing the RSVP LSP-PE-6-PE-1 to the other routers, the LSP shortcut is not used by PE-4 because the metric to reach PE-6 (10) plus the metric of the LSP shortcut from PE-6 to PE-1 (metric 30) is greater than 20.

```
*A:PE-4# show router route-table  
-----  
Route Table (Router: Base)  
-----  
Dest Prefix                               Type   Proto   Age           Pref  
  Next Hop[Interface Name]                 Metric
```

```

-----
192.0.2.1/32                               Remote  ISIS    05h55m29s  15
      192.168.24.1                          20
192.0.2.2/32                               Remote  ISIS    02d04h33m  15
      192.168.24.1                          10
192.0.2.3/32                               Remote  ISIS    01d07h09m  15
      192.168.24.1                          20
192.0.2.4/32                               Local   Local   02d04h37m  0
      system                                0
192.0.2.5/32                               Remote  ISIS    01d12h50m  15
      192.168.45.2                          10
192.0.2.6/32                               Remote  ISIS    05h51m47s  15
      192.168.46.2                          10
192.168.12.0/30                            Remote  ISIS    02d04h33m  15
      192.168.24.1                          20
192.168.13.0/30                            Remote  ISIS    05h55m29s  15
      192.168.24.1                          30
192.168.23.0/30                            Remote  ISIS    02d04h33m  15
      192.168.24.1                          20
192.168.24.0/30                            Local   Local   02d04h37m  0
      int-PE-4-PE-2                        0
192.168.35.0/30                            Remote  ISIS    01d12h50m  15
      192.168.45.2                          20
192.168.45.0/30                            Local   Local   01d12h50m  0
      int-PE-4-PE-5                        0
192.168.46.0/30                            Local   Local   02d04h37m  0
      int-PE-4-PE-6                        0
192.168.56.0/30                            Remote  ISIS    01d07h09m  15
      192.168.45.2                          20
-----
No. of Routes: 14
=====

```

If the metric of the LSP LSP-PE-6-PE-1 is modified to a value between 1 and 9, there is a better metric (less than 20) so that PE-4 will change the next hop via PE-6. First the metric of the LSP is modified to 9:

```

*A:PE-6#config router mpls
  lsp "LSP-PE-6-PE-1" metric 9

```

And checking PE-4's GRT the next hop to reach PE-1 has changed, from next-hop PE-2 to next-hop PE-6 (hence, using the LSP shortcut), and the metric is 19 (10 to reach PE-6 plus metric 9 of the LSP PE-6-PE-1 shortcut):

```

*A:PE-4# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type   Proto   Age           Pref
      Next Hop[Interface Name]                               Metric
-----
192.0.2.1/32                               Remote  ISIS    00h00m10s    15
      192.168.46.2                               19
* Truncated info
-----

```

RSVP Shortcut for IGP Route Resolution

No. of Routes: 14

Because the metric of the LSP shortcut was modified to a value of 9, when displaying the GRT of PE-6 it is noted that the next hops of several routes have changed and are also using the shortcut LSP PE-6-PE-1 because the metric is better than the regular IS-IS metric. It is important to emphasize that IGP shortcuts will not be used to resolve prefixes downstream of the LSP endpoint when the LSP metric is higher than the underlying IGP cumulative metric.

```
*A:PE-6# show router route-table
```

```
Route Table (Router: Base)
```

Dest Prefix	Next Hop[Interface Name]	Type	Proto	Age	Pref
172.16.6.1/32	loopback	Local	Local	01d03h55m	0
192.0.2.1/32	192.0.2.1 (tunneled:RSVP:1)	Remote	ISIS	00h02m58s	15
192.0.2.2/32	192.0.2.1 (tunneled:RSVP:1)	Remote	ISIS	00h02m58s	15
192.0.2.3/32	192.168.56.1	Remote	ISIS	06h11m33s	15
192.0.2.4/32	192.168.46.1	Remote	ISIS	06h11m33s	15
192.0.2.5/32	192.168.56.1	Remote	ISIS	06h11m33s	15
192.0.2.6/32	system	Local	Local	08d04h27m	0
192.168.12.0/30	192.0.2.1 (tunneled:RSVP:1)	Remote	ISIS	00h02m58s	15
192.168.13.0/30	192.0.2.1 (tunneled:RSVP:1)	Remote	ISIS	00h02m58s	15
192.168.23.0/30	192.0.2.1 (tunneled:RSVP:1)	Remote	ISIS	00h02m58s	15
192.168.24.0/30	192.168.46.1	Remote	ISIS	06h11m33s	15
192.168.35.0/30	192.168.56.1	Remote	ISIS	06h11m33s	15
192.168.45.0/30	192.168.46.1	Remote	ISIS	06h11m33s	15
192.168.46.0/30	int-PE-6-PE-4	Local	Local	03d15h03m	0
192.168.56.0/30	int-PE-6-PE-5	Local	Local	08d04h27m	0

No. of Routes: 15

Note that there are also cases where an LDP FEC can resolve to an RSVP LSP, if the user enables the LDP-over-RSVP feature or IGP shortcut feature when prefer-tunnel-in-tunnel is enabled in LDP and the endpoint of the RSVP LSP matches the FEC prefix. For those cases, the metric to the prefix is the sum of the RSVP LSP metric + remaining IGP path cost.

Table 5 provides a summary of the outcome when configuring the forwarding adjacency, LDPoRSVP and RSVP shortcut options at both the IGP instance level and at the LSP level.

Table 5: RSVP LSP Role As Outcome of LSP Level and IGP Level Configuration Options

LSP Level Configuration	IGP Instance Level Configurations					
	advertise-tunnel-link enabled/ rsvp-shortcut enabled/ ldp-over-rsvp enabled	advertise-tunnel-link enabled/ rsvp-shortcut enabled/ ldp-over-rsvp disabled	advertise-tunnel-link enabled/ rsvp-shortcut disabled/ ldp-over-rsvp disabled	advertise-tunnel-link disabled/ rsvp-shortcut disabled/ ldp-over-rsvp disabled	advertise-tunnel-link disabled/ rsvp-shortcut enabled/ ldp-over-rsvp enabled	advertise-tunnel-link disabled/ rsvp-shortcut disabled/ ldp-over-rsvp enabled
igp-shortcut enabled/ldp-over-rsvp enabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	LDP-over-RSVP
igp-shortcut enabled/ldp-over-rsvp disabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	None
igp-shortcut disabled/ldp-over-rsvp enabled	None	None	None	None	None	LDP-over-RSVP
igp-shortcut disabled/ldp-over-rsvp disabled	None	None	None	None	None	None

LSP Relative Metric

It is possible to use relative metrics for IGP shortcuts as per RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*, with the following command:

```
config router mpls lsp <lsp-name> igp-shortcut relative-metric <offset>
```

When this feature is enabled, IGP applies the shortest IGP cost between the endpoints of the LSP, plus the value of a configured offset when computing the cost of the prefix that is resolved to the LSP.

The offset value is optional and can have a value between -10 and 10, and defaults to zero (0). An offset value of zero (0) is used when the relative-metric option is enabled without specifying the offset parameter value. The minimum net cost for the prefix is capped to the value of one (1) after applying the offset:

Prefix cost = max (1, IGP Cost + relative metric offset)

The **relative-metric** option is ignored when advertise-tunnel-link is enabled in IS-IS or OSPF, in that case, the IGP advertises the LSP as a P2P unnumbered link using the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. An LSP with **relative-metric** option enabled cannot be included in the LFA SPF and vice-versa when the **rsvp-shortcut** option is enabled in the IGP (see [LDP/IP FRR LFA for IGP Shortcut Using IS-IS/OSPF on page 1011](#) for more information).

The offset can be used to enforce the preference of the shortcut path over the other paths for the prefix. Using an example, a new CSPF LSP with empty path and relative metric of -10 is created between PE-6 and PE-1. While the operational or absolute metric is 30 (IGP cost and populated in the Tunnel Table Manager, TTM), the metric that the RTM shows is 20 after applying the offset:

```
A:PE-6>config>router>mpls>lsp LSP-PE-6-PE-1-2# info
```

```
-----
to 192.0.2.1
cspf
igp-shortcut relative-metric -10
primary "p"
exit
no shutdown
```

```
A:PE-6# show router tunnel-table 192.0.2.1
```

```
-----
Tunnel Table (Router: Base)
```

```
-----
Destination          Owner Encap TunnelId  Pref    Nexthop      Metric
-----
192.0.2.1/32         rsvp  MPLS   4         7      192.168.56.1  30
-----
```

```
A:PE-6# show router route-table 192.0.2.1
```

```

=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                                Type   Proto   Age      Pref
  Next Hop[Interface Name]                        Metric
-----
192.0.2.1/32                                       Remote ISIS   00h00m23s 15
  192.0.2.1 (tunneled:RSVP:4)                       20
-----
No. of Routes: 1

```

LDP/IP FRR LFA for IGP Shortcut Using IS-IS/OSPF

MPLS LDP/IP FRR LFA for IGP shortcuts adds the use of RSVP-LSP-based IGP shortcuts as a Loop-Free Alternate (LFA) backup to expand the coverage of the IP Fast-Reroute (FRR) capability and the LDP FRR capability for IS-IS and OSPF prefixes. For a detailed description about IP and LDP FRR, refer to [MPLS LDP FRR using ISIS as IGP on page 1101](#).

When an RSVP LSP is used as a shortcut by IS-IS or OSPF, it is included by the SPF as a P2P link and it can also be optionally advertised into the rest of the network by the IGP.

Two LSP-level configuration options are provided:

- The `lfa-protect` option includes the RSVP LSP in both the main SPF and the LFA SPFs. If the prefix primary Next-Hop (NH) is tunneled, no LFA NH is computed. The protection in this case is provided by RSVP FRR. If the prefix primary NH is direct, then an LFA NH is computed. A direct LFA NH is preferred over a tunneled LFA NH. Within each LFA NH type, node-protection is preferred over link-protection. The configuration command is:

```
configure router mpls lsp <lsp-name> igp-shortcut lfa-protect
```

- The `lfa-only` option includes the LSP in the LFA SPFs only so that the introduction of IGP shortcuts does not impact the main SPF decision. The prefix primary NH is always direct and the prefix LFA NH is computed. A direct LFA NH is preferred over a tunneled LFA NH. Within each LFA NH type, node-protection is preferred over link-protection. The configuration command is:

```
configure router mpls lsp <lsp-name> igp-shortcut lfa-only
```

LDP/IP FRR is a local decision so it can be enabled per node and there are no interoperability issues with other nodes. In the topology, PE-2 is provisioned with IS-IS LFA (OSPF configuration for the rest of this section is similar):

```
A:PE-2# configure router isis loopfree-alternate
```

RSVP Shortcut for IGP Route Resolution

The second item to configure is whether LDP or IP FRR is provisioned. To configure ip-fast-reroute, the command is:

```
A:PE-2# configure router ip-fast-reroute
```

Note: IP FRR feature for IS-IS/OSPF prefixes is supported on 7950 XRS, on 7750 SR-7/12/12e in chassis mode D, on the 7450 ESS-6/6v/7/12 in chassis mode D with or without mixed-mode, and 7750 SR-c4/c12.

To configure LDP FRR (no chassis dependency), this command is used:

```
A:PE-2# configure router ldp fast-reroute
```

Note: Although not shown, it is recommended to enable IGP-LDP synchronization per interface to avoid possible traffic blackholes.

Once LFA is enabled in all routers of the topology, looking at PE-2 (the configuration is done only on this node), the following command shows the LFA coverage where 4/5 nodes and 7/10 IPv4 prefixes are protected (IPv6 is not configured). Note that the output shows L1 and L2 because this node is provisioned as an L1-L2 IS-IS router. PE-2, PE-3, PE-4 and PE-5 share the same results, whereas only PE-1 and PE-6 have a 100% of coverage.

```
A:PE-2# show router isis lfa-coverage
```

```
=====
LFA Coverage
=====
```

Topology	Level	Node	IPv4	IPv6
IPv4 Unicast	L1	4/5 (80%)	7/10 (70%)	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	4/5 (80%)	7/10 (70%)	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>config>router# show router isis lfa-coverage
```

```
=====
LFA Coverage
=====
```

Topology	Level	Node	IPv4	IPv6
IPv4 Unicast	L1	5/5 (100%)	11/11 (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	5/5 (100%)	11/11 (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%)

```
=====
```


=====

Taking a deeper look into the IS-IS LFA on PE-2, it can be seen that the node which is not protected is PE-4 (system address 192.0.2.4, since it is the one missing):

```
A:PE-2# show router route-table alternative | match LFA pre-lines 2
192.0.2.1/32          Remote  ISIS      00h44m46s  15
    192.168.12.1          10
    192.168.23.2 (LFA)    20
192.0.2.3/32          Remote  ISIS      00h44m46s  15
    192.168.23.2          10
    192.168.12.1 (LFA)    20
192.0.2.5/32          Remote  ISIS      00h44m46s  15
    192.168.23.2          20
    192.168.24.2 (LFA)    20
192.0.2.6/32          Remote  ISIS      00h44m46s  15
    192.168.24.2          20
    192.168.23.2 (LFA)    30
192.168.13.0/30       Remote  ISIS      00h44m46s  15
    192.168.12.1          20
    192.168.23.2 (LFA)    30
192.168.35.0/30       Remote  ISIS      00h44m46s  15
    192.168.23.2          20
    192.168.12.1 (LFA)    30
192.168.56.0/30       Remote  ISIS      00h44m46s  15
    192.168.23.2          30
    192.168.24.2 (LFA)    30
```

LFA is improved by taking advantage of RSVP shortcuts when it is properly provisioned. The reason why PE-4 cannot be protected with an LFA path is because the direct NH is using the direct link between PE-2 and PE-4 (the shortest IGP) and the intended LFA path through PE-3 is not valid (when LFA tries to find an alternate path via PE-3, the IGP cost from PE-3 to PE-4 is the same going via PE-5 then the path back via PE-2, invalidating that LFA calculation as there is a loop). This is normal as PE-2, PE-3, PE-4 and PE-5 are forming a ring. LFA coverage is increased by adding a link between PE-2 and PE-5, which can be done using a physical link or a virtual link with an RSVP shortcut. From the two possible options (lfa-only and lfa-protect), a new LSP “LSP-PE-2-PE-5” is configured with igp-shortcut lfa-only.

```
A:PE-2>config>router>mpls# info
-----
    path "to-PE-5-LFA"
        hop 1 192.0.2.3 strict
        hop 2 192.0.2.5 strict
        no shutdown
    exit
    lsp "LSP-PE-2-PE-5"
        to 192.0.2.5
        igp-shortcut lfa-only
        primary "to-PE-5-LFA"
        exit
        no shutdown
    exit
```

RSVP Shortcut for IGP Route Resolution

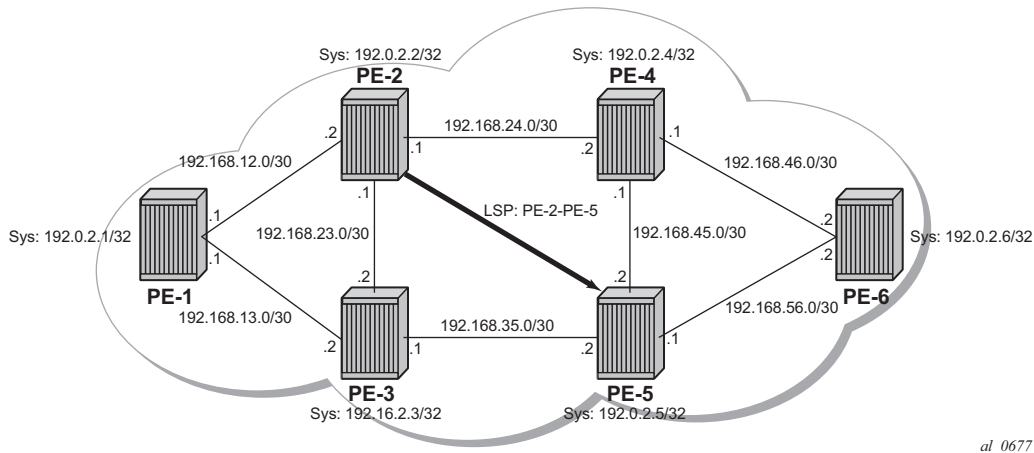


Figure 144: RSVP Shortcuts LFA Use Case Example

Now the coverage is 100% as shown by the output:

```
A:PE-2# show router isis lfa-coverage
=====
LFA Coverage
=====
Topology          Level  Node          IPv4           IPv6
-----
IPV4 Unicast     L1     5/5 (100%)    10/10 (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     5/5 (100%)    10/10 (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%)
=====
```

The GRT details the prefix information after the new LFA calculation using the lfa-only option (the shortcut is used by LFA SPF). Note that the metric from PE-2 to PE-4 is the maximum plus the IGP cost (16777215 + 10) and that the shortcut is also used to protect the rest of the previously unprotected prefixes:

```
*A:PE-2# show router route-table alternative | match LFA pre-lines 2
192.0.2.1/32          Remote  ISIS        01h22m11s  15
    192.168.12.1          10
    192.168.23.2 (LFA)    20
192.0.2.3/32          Remote  ISIS        01h22m11s  15
    192.168.23.2          10
    192.168.12.1 (LFA)    20
192.0.2.4/32          Remote  ISIS        01h22m11s  15
    192.168.24.2          10
    192.0.2.5 (LFA) (tunneled:RSVP:5) 16777225
```

192.0.2.5/32	Remote	ISIS	00h04m52s	15
192.168.23.2			20	
192.168.24.2 (LFA)			20	
192.0.2.6/32	Remote	ISIS	01h22m11s	15
192.168.24.2			20	
192.168.23.2 (LFA)			30	
192.168.13.0/30	Remote	ISIS	01h22m11s	15
192.168.12.1			20	
192.168.23.2 (LFA)			30	
192.168.35.0/30	Remote	ISIS	00h04m52s	15
192.168.23.2			20	
192.168.12.1 (LFA)			30	
192.168.45.0/30	Remote	ISIS	01h22m11s	15
192.168.24.2			20	
192.0.2.5 (LFA) (tunneled:RSVP:5)			16777235	
192.168.46.0/30	Remote	ISIS	01h22m11s	15
192.168.24.2			20	
192.0.2.5 (LFA) (tunneled:RSVP:5)			16777235	
192.168.56.0/30	Remote	ISIS	00h04m52s	15
192.168.23.2			30	
192.168.24.2 (LFA)			30	

The tunnel table shows the RSVP LSP used as a shortcut and its operational metric.

```
A:PE-2# show router tunnel-table 192.0.2.5
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId Pref  NextHop      Metric
-----
192.0.2.5/32         rsvp  MPLS    5       7    192.168.23.2 16777215
=====
```

Now, if the LSP “LSP-PE-2-PE-5” is provisioned with lfa-protect instead of lfa-only, the result is that the LSP "LSP-PE-2-PE-5" is used by normal SPF to define the primary NH and it is not used by LFA SPF anymore.

```
A:PE-2# configure router mpls lsp "LSP-PE-2-PE-5" igp-shortcut lfa-protect
```

The coverage when lfa-protect is used also shows a 100% for nodes. The 112% coverage for prefixes shown below is not correct and should be 100% since the IGP shortcut used as a primary next-hop to reach PE-5 from PE-2 is providing protection for prefixes reachable via PE-5. The display issue will be fixed in a later release and will be referenced by [201872] in the release notes.

```
A:PE-2# show router isis lfa-coverage
=====
LFA Coverage
=====
Topology          Level  Node          IPv4          IPv6
-----
IPV4 Unicast     L1     4/4 (100%)    9/8 (112%)   0/0 (0%)
IPV6 Unicast     L1     0/0 (0%)     0/0 (0%)     0/0 (0%)
```

RSVP Shortcut for IGP Route Resolution

```

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%) 9/8 (112%) 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%)

```

In this case the GRT looks as follows, the main difference being that now PE-5 (192.0.2.5) has a direct shortcut from PE-2:

```
A:PE-2# 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	Remote	ISIS	01h32m11s	15
192.168.12.1			10	
192.168.23.2 (LFA)			20	
192.0.2.2/32	Local	Local	05h05m04s	0
system			0	
192.0.2.3/32	Remote	ISIS	01h32m11s	15
192.168.23.2			10	
192.168.12.1 (LFA)			20	
192.0.2.4/32	Remote	ISIS	01h32m11s	15
192.168.24.2			10	
192.0.2.5 (LFA) (tunneled:RSVP:5)			16777225	
192.0.2.5/32	Remote	ISIS	00h03m07s	15
192.0.2.5 (tunneled:RSVP:5)			16777215	
192.0.2.6/32	Remote	ISIS	01h32m11s	15
192.168.24.2			20	
192.168.23.2 (LFA)			30	
192.168.12.0/30	Local	Local	05h04m51s	0
int-PE-2-PE-1			0	
192.168.13.0/30	Remote	ISIS	01h32m11s	15
192.168.12.1			20	
192.168.23.2 (LFA)			30	
192.168.23.0/30	Local	Local	05h04m51s	0
int-PE-2-PE-3			0	
192.168.24.0/30	Local	Local	05h04m51s	0
int-PE-2-PE-4			0	
192.168.35.0/30	Remote	ISIS	00h03m07s	15
192.168.23.2			20	
192.168.12.1 (LFA)			30	
192.168.45.0/30	Remote	ISIS	01h32m11s	15
192.168.24.2			20	
192.0.2.5 (LFA) (tunneled:RSVP:5)			16777235	
192.168.46.0/30	Remote	ISIS	01h32m28s	15
192.168.24.2			20	
192.0.2.5 (LFA) (tunneled:RSVP:5)			16777235	
192.168.56.0/30	Remote	ISIS	00h03m24s	15
192.168.24.2			30	
192.168.23.2 (LFA)			40	

RSVP Shortcut for IGP Route Resolution

```

-----
172.16.6.1/32                               Local   Local   01d16h00m  0
    loopback                                0
192.0.2.1/32                                Remote  LDP     00h00m11s  9
    192.168.56.1 (tunneled)                 30
192.0.2.2/32                                Remote  LDP     00h00m11s  9
    192.168.56.1 (tunneled)                 30
192.0.2.3/32                                Remote  LDP     00h00m11s  9
    192.168.56.1 (tunneled)                 20
192.0.2.5/32                                Remote  LDP     00h23m35s  9
    192.168.56.1 (tunneled)                 10
192.0.2.6/32                                Local   Local   08d16h32m  0
    system
192.168.12.0/30                             Remote  ISIS    00h00m11s  15
    192.168.56.1                             40
192.168.13.0/30                             Remote  ISIS    00h00m11s  15
    192.168.56.1                             30
192.168.23.0/30                             Remote  ISIS    00h00m11s  15
    192.168.56.1                             30
192.168.35.0/30                             Remote  ISIS    18h16m26s  15
    192.168.56.1                             20
192.168.56.0/30                             Local   Local   08d16h31m  0
    int-PE-6-PE-5                             0
-----
No. of Routes: 11
=====

```

The Tunnel Table shows the LSPs available for the shortcuts, and hence these are used in the GRT for LDP (but not for RSVP):

```

*A:PE-6>config>router# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination      Owner Encap TunnelId  Pref    Nexthop      Metric
-----
192.0.2.1/32     ldp   MPLS   -        9       192.168.56.1  30
192.0.2.2/32     ldp   MPLS   -        9       192.168.56.1  30
192.0.2.3/32     rsvp  MPLS   2        7       192.168.56.1  20
192.0.2.3/32     ldp   MPLS   -        9       192.168.56.1  20
192.0.2.5/32     ldp   MPLS   -        9       192.168.56.1  10
=====

```

So far, LDP shortcuts are preferred over the IGP next-hops for the system addresses (router-id). After enabling RSVP shortcuts under the IS-IS context (**config router isis rsvp-shortcut**), the changes in the GRT are:

```

*A:PE-6>config>router# show router route-table next-hop-type tunneled
=====
Route Table (Router: Base)
=====
Dest Prefix      Type   Proto   Age      Pref
  Next Hop[Interface Name]           Metric
-----
192.0.2.1/32     Remote ISIS   00h00m03s  15

```

```

192.0.2.3 (tunneled:RSVP:2) 30
192.0.2.2/32 Remote ISIS 00h00m03s 15
192.0.2.3 (tunneled:RSVP:2) 30
192.0.2.3/32 Remote ISIS 00h00m03s 15
192.0.2.3 (tunneled:RSVP:2) 20
192.0.2.5/32 Remote LDP 00h25m35s 9
192.168.56.1 (tunneled) 10
192.168.12.0/30 Remote ISIS 00h00m03s 15
192.0.2.3 (tunneled:RSVP:2) 40
192.168.13.0/30 Remote ISIS 00h00m03s 15
192.0.2.3 (tunneled:RSVP:2) 30
192.168.23.0/30 Remote ISIS 00h00m03s 15
192.0.2.3 (tunneled:RSVP:2) 30

```

No. of Routes: 7

The GRT shows that PE-6 is using an LDP shortcut to reach PE-5, but PE-6 is using the RSVP shortcut to reach not only PE-3 system address, but also PE-1 and PE-2 routes (including all interfaces) which were behind the RSVP LSP shortcut.

In summary, the behavior is:

- When resolving a prefix, SPF picks the RSVP shortcut next-hop if there is an RSVP LSP directly to that address regardless of the IGP path cost compared to the IGP next-hop. When multiple RSVP LSPs to that address exist and all have the same lowest metric, if ECMP is enabled on the system, the LSP with lowest tunnel ID is chosen. In this example, if LSP “LSP-PE-6-PE-3” is provisioned with a metric of 100 (IGP metric is 20), the GRT shows that the PE-3 system address is reachable via the LSP.

```

*A:PE-6# show router route-table 192.0.2.3
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.3/32                      Remote ISIS   00h02m51s  15
  192.0.2.3 (tunneled:RSVP:2)                100
-----
No. of Routes: 1

```

- SPF also picks the RSVP LSP shortcut if both the LSP path and the IGP path to the prefix are via the tail-end of the LSP. This is regardless of the path cost compared to the IGP next-hop. When paths over multiple RSVP shortcuts have the same lowest cost, if ECMP is enabled on the system, the LSP with lowest tunnel ID is chosen. In this example, 192.168.13.0 and 192.168.23.0 are using the shortcut but 192.168.12.0 is not.

```

*A:PE-6# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----

```

RSVP Shortcut for IGP Route Resolution

```
-----  
192.168.12.0/30                               Remote  ISIS    00h00m04s  15  
192.168.56.1                                  40  
192.168.13.0/30                               Remote  ISIS    00h10m56s  15  
192.0.2.3 (tunneled:RSVP:2)                  110  
192.168.23.0/30                               Remote  ISIS    00h00m04s  15  
192.0.2.3 (tunneled:RSVP:2)                  110
```


LDP/RSVP LSP Shortcut for BGP NH Resolution

Using LDP/RSVP LSP shortcuts for resolving BGP next-hops allows IPv4 packet forwarding to routes resolved via a BGP next-hop using an LDP/RSVP LSP instead of using a regular IP next-hop. In the network topology of Figure 3, both PE-3 and PE-6 have a single peer configured, initially without any shortcuts enabled under the BGP context. Also, one static route is configured in PE-3 and PE-6 and that is redistributed into BGP. The relevant configuration on PE-3 is the following:

```
*A:PE-3#config router
  interface "static-route"
    address 172.16.33.1/30
    port 1/1/3:33
  exit
  autonomous-system 65536

  static-route 10.10.10.0/24 next-hop 172.16.33.2

  policy-options
    begin
    policy-statement "static-routes"
      description "export static-routes for I-BGP"
      entry 10
        from
          protocol static
        exit
        to
          protocol bgp
        exit
        action accept
          next-hop-self
        exit
      exit
    exit
  commit
exit

  bgp
    export "static-routes"
    group "ibgp"
      type internal
      neighbor 192.0.2.6
    exit
  exit
exit
```

LDP/RSVP LSP Shortcut for BGP NH Resolution

Checking the static route received on PE-6 via BGP, the next-hop is the PE-3 system address:

```
*A:PE-6>config>router# show router bgp routes 10.10.10.0/24 detail
=====
BGP Router ID:192.0.2.6      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
Original Attributes

Network       : 10.10.10.0/24
Nexthop      : 192.0.2.3
Path Id      : None
From         : 192.0.2.3
Res. Nexthop : 192.168.56.1
Local Pref.    : 100                               Interface Name : int-PE-6-PE-5
Aggregator AS  : None                               Aggregator    : None
Atomic Aggr.  : Not Atomic                           MED           : None
AIGP Metric    : None
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                               Peer Router Id : 192.0.2.3
Fwd Class      : None                               Priority       : None
Flags          : Used Valid Best Incomplete
Route Source   : Internal
AS-Path        : No As-Path
Neighbor-AS    : N/A

Modified Attributes

Network       : 10.10.10.0/24
Nexthop      : 192.0.2.3
Path Id      : None
From         : 192.0.2.3
Res. Nexthop : 192.168.56.1
* Truncated info
-----
-----
Routes : 1
```

The BGP peering configuration possibilities are LDP, RSVP, or MPLS (it chooses RSVP if available, if not LDP), and also disabling the IGP is allowed (meaning that unless there is a shortcut, the BGP peering will not fall back to IGP):

```
*A:PE-6>config>router>bgp# igp-shortcut
- igp-shortcut {ldp|rsvp-te|mpls} [disallow-igp]
- no igp-shortcut

<ldp|rsvp-te|mpls> : keywords
<disallow-igp>    : keyword
```

When enabling LDP shortcuts (**config>router>bgp>igp-shortcut ldp**) on PE-6, the output changes showing the detail of the received BGP route indicating that the next hop is resolved using LDP:

```
*A:PE-6>config>router# show router bgp routes 10.10.10.0/24 detail
=====
BGP Router ID:192.0.2.6      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
Original Attributes
Network       : 10.10.10.0/24
Nexthop      : 192.0.2.3
Path Id      : None
From         : 192.0.2.3
Res. Nexthop : 192.168.56.1 (LDP)
Local Pref.    : 100
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
AIGP Metric    : None
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None
Fwd Class      : None
Flags          : Used Valid Best Incomplete
Route Source   : Internal
AS-Path        : No As-Path
Neighbor-AS    : N/A
Peer Router Id : 192.0.2.3
Priority        : None
Interface Name : int-PE-6-PE-5
Aggregator     : None
MED            : None

Modified Attributes

Network       : 10.10.10.0/24
Nexthop       : 192.0.2.3
Path Id       : None
From          : 192.0.2.3
Res. Nexthop  : 192.168.56.1 (LDP)
* Truncated info
```

LDP/RSVP LSP Shortcut for BGP NH Resolution

```
-----  
-----  
Routes : 1
```

The GRT output command also shows that the route is reachable using LDP (indicated as tunneled):

```
*A:PE-6# show router route-table next-hop-type tunneled  
-----  
Route Table (Router: Base)  
-----  
Dest Prefix                               Type   Proto   Age           Pref  
  Next Hop[Interface Name]                Metric  
-----  
10.10.10.0/24                             Remote BGP     00h07m51s   170  
  192.0.2.3 (tunneled)                    0  
-----  
No. of Routes: 1  
-----
```

An RSVP LSP is then configured between PE-6 and PE-3 as follows.

```
*A:PE-6#config router mpls  
-----  
no shortcut-local-ttl-propagate  
no shortcut-transit-ttl-propagate  
interface "system"  
exit  
interface "int-PE-6-PE-5"  
exit  
interface "int-PE-6-PE-4"  
exit  
path "p"  
no shutdown  
exit  
lsp "LSP-PE-6-PE-3"  
to 192.0.2.3  
cspf  
primary "p"  
exit  
no shutdown  
exit  
no shutdown
```

It can be seen that the LSP LSP-PE-6-PE-3 is up and running:

```
*A:PE-6>config>router# show router mpls lsp "LSP-PE-6-PE-3" path detail  
LSP LSP-PE-6-PE-3 Path p  
-----  
LSP Name      : LSP-PE-6-PE-3                Path LSP ID : 49190  
From          : 192.0.2.6                    To          : 192.0.2.3  
Adm State     : Up                          Oper State  : Up  
Path Name     : p                           Path Type   : Primary  
Path Admin    : Up                          Path Oper   : Up
```

```

OutInterface: 1/1/4:1
Path Up Time: 0d 00:03:32
Retry Limit : 0
RetryAttempt: 0
SetupPriorit*: 7
Preference : n/a
Bandwidth : No Reservation
Hop Limit : 255
Backup CT : None
MainCT Retry: n/a
  Rem :
Oper CT : 0
Record Route: Record
Oper MTU : 9194
Adaptive : Enabled
Include Grps:
None
Path Trans : 39
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  192.168.56.2(192.0.2.6)
  -> 192.168.56.1(192.0.2.5)
  -> 192.168.35.1(192.0.2.3)
ComputedHops:
  192.168.56.2 -> 192.168.56.1 -> 192.168.35.1
ResigEligib*: False
LastResignal: n/a

Out Label : 131067
Path Dn Time: 0d 00:00:00
Retry Timer : 30 sec
NextRetryIn : 0 sec
Hold Priorit*: 0

Oper Bw : 0 Mbps
Class Type : 0

MainCT Retry: 0
  Limit :

Record Label: Record
Neg MTU : 9194
Oper Metric : 20
Exclude Grps:
None
CSPF Queries: 3
Failure Node: n/a

Record Label : N/A
Record Label : 131067
Record Label : 131070

CSPF Metric : 20

```

After enabling **config>router>bgp>igp-shortcut mpls**, the output shows that the BGP peer is reachable using an RSVP LSP (switched from LDP to RSVP since RSVP is preferred):

```

*A:PE-6# show router bgp routes ipv4 10.10.10.0/24 detail
=====
BGP Router ID:192.0.2.6      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
Original Attributes

Network      : 10.10.10.0/24
Nexthop     : 192.0.2.3
Path Id     : None
From        : 192.0.2.3
Res. Nexthop : 192.168.46.1 (RSVP LSP: 2)
Local Pref. : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector   : None
Community   : No Community Members

Interface Name : int-PE-6-PE-4
Aggregator     : None
MED            : None

```

LDP/RSVP LSP Shortcut for BGP NH Resolution

```
Cluster      : No Cluster Members
Originator Id : None                Peer Router Id : 192.0.2.3
Fwd Class    : None                Priority       : None
Flags        : Used Valid Best Incomplete
Route Source : Internal
AS-Path      : No As-Path
Neighbor-AS  : N/A
```

Modified Attributes

```
Network      : 10.10.10.0/24
Nexthop      : 192.0.2.3
Path Id      : None
From         : 192.0.2.3
Res. Nexthop : 192.168.46.1 (RSVP LSP: 2)
* Truncated info
```

```
-----
Routes : 1
```

The GRT output command also shows that the route is reachable using RSVP (indicated as tunneled:RSVP:2):

```
*A:PE-6>config>router>bgp# show router route-table next-hop-type tunneled
=====
Route Table (Router: Base)
=====
Dest Prefix          Type   Proto   Age           Pref
  Next Hop[Interface Name]           Metric
-----
10.10.10.0/24       Remote BGP     00h00m10s  170
  192.0.2.3 (tunneled:RSVP:2)           0
-----
No. of Routes: 1
=====
```

If the RSVP LSP is shutdown (*A:PE-6>config>router# mpls lsp "LSP-PE-6-PE-3" shutdown) with “config router bgp igp-shortcut mpls” enabled, the system reverts back to the LDP LSP:

```
*A:PE-6# show router bgp routes 10.10.10.0/24 detail
=====
BGP Router ID:192.0.2.6      AS:65000      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP IPv4 Routes
=====
Original Attributes

Network      : 10.10.10.0/24
Nexthop      : 192.0.2.3
From         : 192.0.2.3
```

```

Res. Nexthop   : 192.168.56.1 (LDP)
Local Pref.    : 100                               Interface Name : int-PE-6-PE-5
Aggregator AS : None                               Aggregator    : None
Atomic Aggr.  : Not Atomic                         MED           : None
Community     : No Community Members
Cluster       : No Cluster Members
Originator Id : None                               Peer Router Id : 192.0.2.3
Flags         : Used Valid Best Incomplete
AS-Path       : No As-Path

```

Modified Attributes

```

Network       : 10.10.10.0/24
Nexthop       : 192.0.2.3
From          : 192.0.2.3
Res. Nexthop  : 192.168.56.1 (LDP)
* Truncated info

```

```

-----
Routes : 1

```

When `igp-shortcut` is enabled at the BGP level, all RSVP LSPs originating on this node are eligible to be used by default as long as the destination address of the LSP is configured in `config>router>mpls>lsp` and corresponds to that of the BGP next-hop for that prefix. It is also possible to exclude a specific RSVP LSP from BGP next-hop resolution, similar to the exclusion of a specific RSVP LSP being used as a shortcut for resolving IGP routes. In this example, if the RSVP LSP "LSP-PE-6-PE-3" is excluded to be eligible for BGP next-hop resolution, it reverts back to LDP.

```

*A:PE-6# configure router mpls lsp "LSP-PE-6-PE-3" no bgp-shortcut
*A:PE-6# show router route-table 10.10.10.0
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                       Metric
-----
10.10.10.0/24                                     Remote BGP     00h17m09s  170
  192.0.2.3 (tunneled)                            0
-----
No. of Routes: 1

```

If the configuration is using `config>router>bgp>igp-shortcut mpls disallow-igp`, and neither LDP nor RSVP LSPs are available, the remote route received via BGP is removed from the GRT although the BGP peer session remains up. A field in the detailed show BGP route output indicates that the next hop is "Unresolved":

```

*A:PE-6# show router bgp routes 10.10.10.0/24 detail
=====
BGP Router ID:192.0.2.6      AS:65000      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid

```

LDP/RSVP LSP Shortcut for BGP NH Resolution

Origin codes : i - IGP, e - EGP, ? - incomplete, > - best

=====
 BGP IPv4 Routes
 =====

 Original Attributes

Network : 10.10.10.0/24
Nexthop : 192.0.2.3
From : 192.0.2.3
Res. Nexthop : Unresolved
 Local Pref. : 100 Interface Name : int-PE-6-PE-5
 Aggregator AS : None Aggregator : None
 Atomic Aggr. : Not Atomic MED : None
 Community : No Community Members
 Cluster : No Cluster Members
 Originator Id : None Peer Router Id : 192.0.2.3
 Flags : Invalid Incomplete Nexthop-Unresolved
 AS-Path : No As-Path

Modified Attributes

Network : 10.10.10.0/24
 Nexthop : 192.0.2.3
 From : 192.0.2.3
 Res. Nexthop : Unresolved
 Local Pref. : 100 Interface Name : int-PE-6-PE-5
 Aggregator AS : None Aggregator : None
 Atomic Aggr. : Not Atomic MED : None
 Community : No Community Members
 Cluster : No Cluster Members
 Originator Id : None Peer Router Id : 192.0.2.3
 Flags : Invalid Incomplete Nexthop-Unresolved
 AS-Path : No As-Path

 Routes : 1

As the route is unresolved, it does not appear in the GRT:

*A:PE-6# show router route-table

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

Dest Prefix	Type	Proto	Age	Pref
Next Hop[Interface Name]			Metric	
20.20.20.0/24	Remote	Static	03d06h45m	5
172.16.66.2			1	
172.16.6.1/32	Local	Local	05d00h03m	0
loopback			0	
172.16.66.0/30	Local	Local	03d06h45m	0
static-route			0	
192.0.2.1/32	Remote	ISIS	03d06h34m	15
192.168.56.1			30	
192.0.2.2/32	Remote	ISIS	03d07h52m	15

192.168.46.1			20	
192.0.2.3/32	Remote	ISIS	03d06h34m	15
192.168.56.1			20	
192.0.2.4/32	Remote	ISIS	03d07h53m	15
192.168.46.1			10	
192.0.2.5/32	Remote	ISIS	04d02h20m	15
192.168.56.1			10	
192.0.2.6/32	Local	Local	12d00h36m	0
system			0	
192.168.12.0/30	Remote	ISIS	03d07h52m	15
192.168.46.1			30	
192.168.13.0/30	Remote	ISIS	03d06h34m	15
192.168.56.1			30	
192.168.23.0/30	Remote	ISIS	03d07h52m	15
192.168.46.1			30	
192.168.24.0/30	Remote	ISIS	03d07h53m	15
192.168.46.1			20	
192.168.35.0/30	Remote	ISIS	04d02h20m	15
192.168.56.1			20	
192.168.45.0/30	Remote	ISIS	03d07h53m	15
192.168.46.1			20	
192.168.46.0/30	Local	Local	03d07h54m	0
int-PE-6-PE-4			0	
192.168.56.0/30	Local	Local	12d00h35m	0
int-PE-6-PE-5			0	

No. of Routes: 17

MPLS/GRE Shortcut for BGP NH Resolution within a VRF

Using RSVP/LDP or GRE shortcuts for resolving BGP next-hops within a VPRN (Virtual Private Routed Network), also known as auto-bind, allows a VPRN service to automatically resolve the BGP next-hop for VPRN routes to an MPLS LSP or a GRE tunnel. Three possible mechanisms to provide transport tunnels for forwarding traffic between PE routers within an RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*, network exist:

- RSVP-TE protocol to create tunnel LSPs between PE routers.
- LDP protocol to create tunnel LSPs between PE routers.
- GRE tunnels between PE routers.

These transport tunneling mechanisms provide the flexibility to use dynamically created LSPs where the service tunnels are automatically bound (the “auto-bind” feature), and the ability to provide certain VPN services with their own transport tunnels by explicitly binding SDPs if desired. All services using the auto-bind feature use the same set of LSPs, which does not allow for alternate tunneling mechanisms (like GRE) or the ability to craft sets of LSPs with bandwidth reservations for specific customers, as is available with explicit SDPs for the service.

The auto-bind configuration is as follows:

```
A:PE-2>config>service>vprn# auto-bind
- auto-bind {ldp|gre|rsvp-te|mpls|mpls-gre}
- no auto-bind

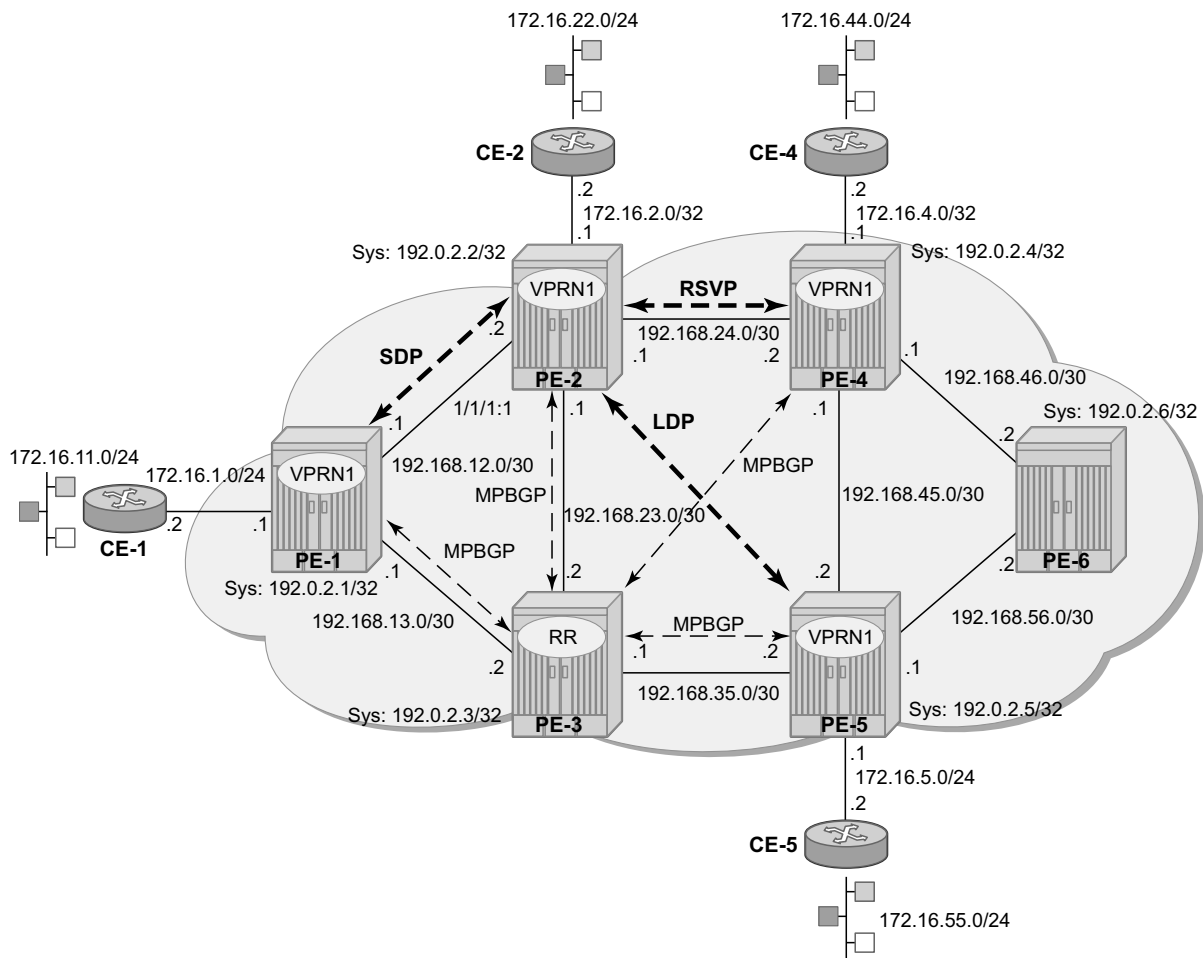
<ldp|gre|rsvp-te|m*> : ldp|gre|rsvp-te|mpls|mpls-gre
```

Parameter descriptions:

- **ldp** — Specifies LDP based LSPs should be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance.
- **gre** — Specifies GRE based tunnels to be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance. GRE is out of the scope regarding shortcuts, refer to SR OS documentation for further details.
- **rsvp-te** — Specifies RSVP-TE LSPs should be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance.
- **mpls** — Chooses an existing RSVP-TE LSP if available, otherwise use LDP.
- **mpls-gre** — When there are multiple tunnels to the BGP next-hop address, the tunnel with the lowest tunnel-table preference value is selected (first RVSP, then LDP, then GRE).

In all cases, if an explicit spoke-sdp is specified in the VPRN, it is always preferred over automatically selected tunnels (even if the SDP is down the route becomes inactive, there is no fallback to the automatic selection).

The network is configured according to the topology shown in Figure 6. Four PEs (PE-1, PE-2, PE-4 and PE-5) are connected forming a meshed IP-VPN (named VPRN 1) with “autobind mpls-gre”, using a route reflector on PE-3 for MP-BGP peering. All PEs have LDP tunnels enabled so at a minimum all can establish LDP shortcut tunnels to the others. In order to have not only LDP but also RSVP-TE LSPs and static SDPs (using an RSVP LSP) in the network, a mix of tunneling methods is configured. For the sake of simplicity, a closer view on PE-2 only, provides all details about the shortcuts created by auto-bind. PE-2 has a static SDP (RSVP based) with PE-1, an RSVP LSP with PE-4, and an LDP LSP with PE-5. Every PE has a CE connected, so each PE has an interface connected to the CE as well as a static route to a CE LAN (although redistribution routing policies are needed, they are not shown for simplicity).



OSSG627

Figure 146: Shortcuts Within a VRF Topology Network

MPLS/GRE Shortcut for BGP NH Resolution within a VRF

On PE-2, the following output shows the configuration of VPRN1:

```
A:PE-2#config service vprn 1 customer 1 create
    vrf-import "VPN1-import"
    vrf-export "VPN1-export"
    route-distinguisher 65002:1
    auto-bind mpls-gre
    interface "to-ce2" create
        address 172.16.2.1/24
        sap 1/1/3:1 create
    exit
exit
static-route 172.16.22.0/24 next-hop 172.16.2.2
spoke-sdp 1 create
exit
no shutdown
```

As previously mentioned, regarding IP-VPN meshed connectivity, the configuration shows that there is a static SDP 1 (pointing to PE-1), and the rest of the configuration is just “auto-bind mpls-gre”. On PE-2, the connectivity towards the other PEs in the network can be verified by checking VPRN 1:

```
A:PE-2# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                               Type  Proto  Age           Pref
  Next Hop[Interface Name]                       Metric
-----
172.16.1.0/24                                     Remote BGP VPN 00h12m06s 170
    192.0.2.1 (tunneled)                          0
172.16.2.0/24                                     Local  Local  02d03h54m  0
    to-ce2                                           0
172.16.4.0/24                                     Remote BGP VPN 00h01m02s 170
    192.0.2.4 (tunneled:RSVP:2)                    0
172.16.5.0/24                                     Remote BGP VPN 00h01m29s 170
    192.0.2.5 (tunneled)                          0
172.16.11.0/24                                    Remote BGP VPN 00h12m06s 170
    192.0.2.1 (tunneled)                          0
172.16.22.0/24                                    Remote Static 02d03h54m  5
    172.16.2.2                                       1
172.16.44.0/24                                    Remote BGP VPN 00h01m02s 170
    192.0.2.4 (tunneled:RSVP:2)                    0
172.16.55.0/24                                    Remote BGP VPN 00h01m29s 170
    192.0.2.5 (tunneled)                          0
-----
No. of Routes: 8
```

As can be seen, there are eight routes since every PE has two routes (one direct PE-CE interface and one static route), so six routes are received from other PEs via MP-BGP. The VPRN 1 routing table can be understood by looking at the tunnel table (active LSPs for remote system-ids):

```
A:PE-2# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination                Owner Encap TunnelId  Pref    Nexthop      Metric
-----
192.0.2.1/32                sdp   MPLS   1         5       192.0.2.1     0
192.0.2.1/32                rsvp  MPLS   1         7       192.168.12.1  10
192.0.2.1/32                ldp   MPLS   -         9       192.168.12.1  10
192.0.2.4/32                rsvp  MPLS   2         7       192.168.24.2  10
192.0.2.4/32                rsvp  MPLS   3         7       192.168.23.2  16777215
192.0.2.4/32                ldp   MPLS   -         9       192.168.24.2  10
192.0.2.5/32                ldp   MPLS   -         9       192.168.23.2  20
* Truncated info
-----
```

The tunnel-table shows one entry per LSP per remote PE. The following tunnel selection rules apply:

- SDP has lowest (best) preference, followed by RSVP then by LDP.
- If the preference is the same, the lowest metric is selected (ECMP is possible with LDP).

PE-2 has three possibilities to reach PE-1 (192.0.2.1): an SDP Tunnel ID 1 with preference 5, an RSVP Tunnel ID 1 with preference 7, and an LDP LSP with preference 9. As SDP Tunnel ID 1 has the lowest preference, it is the chosen option. PE-2 has three possibilities to reach PE-4 (192.0.2.4): an RSVP Tunnel ID 2 with preference 7 and metric 10, an RSVP Tunnel ID 3 with preference 7 and metric 16777215, and an LDP LSP with preference 9; hence RSVP Tunnel ID 2 is selected. PE-2 only has one option to reach PE-5 and PE-6 (192.0.2.5 and .6) using an LDP LSP. As the VPRN 1 output does not show the details of the tunneling, displaying the FIB on router VPRN 1 provides more detailed information:

```
A:PE-2# show router 1 fib 1
=====
FIB Display
=====
Prefix                Protocol
NextHop
-----
172.16.1.0/24          BGP_VPN
  192.0.2.1 (VPRN Label:262142 Transport:SDP:1)
172.16.2.0/24          LOCAL
  172.16.2.0 (to-ce2)
172.16.4.0/24          BGP_VPN
  192.0.2.4 (VPRN Label:131070 Transport:RSVP LSP:2)
172.16.5.0/24          BGP_VPN
  192.0.2.5 (VPRN Label:131070 Transport:LDP)
172.16.11.0/24         BGP_VPN
  192.0.2.1 (VPRN Label:262142 Transport:SDP:1)
172.16.22.0/24         STATIC
  172.16.2.2 (to-ce2)
172.16.44.0/24         BGP_VPN
  192.0.2.4 (VPRN Label:131070 Transport:RSVP LSP:2)
172.16.55.0/24         BGP_VPN
  192.0.2.5 (VPRN Label:131070 Transport:LDP)
```

MPLS/GRE Shortcut for BGP NH Resolution within a VRF

```
-----  
Total Entries : 8  
-----
```

The FIB shows the chosen transport tunnel, specifying SDP ID, RSVP Tunnel ID, and LDP, as well as service label information linked to the routes.

As static SDP tunnels are preferred over dynamic tunnels (RSVP or LDP auto-bind), when the static SDP 1 is shutdown with “config service sdp 1 shutdown” or the LSP goes down (there is no fallback to dynamic tunneling), the associated routes are removed:

```
*A:PE-2# show router 1 fib 1  
=====
```

Prefix	NextHop	Protocol
172.16.2.0/24	172.16.2.0 (to-ce2)	LOCAL
172.16.4.0/24	192.0.2.4 (VPRN Label:131070 Transport:RSVP LSP:2)	BGP_VPN
172.16.5.0/24	192.0.2.5 (VPRN Label:131070 Transport:LDP)	BGP_VPN
172.16.22.0/24	172.16.2.2 (to-ce2)	STATIC
172.16.44.0/24	192.0.2.4 (VPRN Label:131070 Transport:RSVP LSP:2)	BGP_VPN
172.16.55.0/24	192.0.2.5 (VPRN Label:131070 Transport:LDP)	BGP_VPN

```
-----  
Total Entries : 6  
-----
```

To avoid this fallback issue, the configuration is modified and the manual spoke-sdps are removed in the configuration on PE-1 and PE-2, the rest of the configuration remains the same. Now the connectivity between PE-1 and PE-2 is using an RSVP LSP, as shown in the PE-1 output below (RSVP LSP which was used by SDP 1 has disappeared):

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

Dest Prefix[Flags]	Type	Proto	Age	Pref
Next Hop[Interface Name]			Metric	
172.16.1.0/24	Local	Local	07h00m15s	0
to-ce1			0	
172.16.2.0/24	Remote	BGP VPN	00h00m15s	170
192.0.2.2 (tunneled:RSVP:2)			0	
172.16.4.0/24	Remote	BGP VPN	00h09m28s	170
192.0.2.4 (tunneled)			0	
172.16.5.0/24	Remote	BGP VPN	00h09m28s	170
192.0.2.5 (tunneled)			0	

```

172.16.11.0/24          Remote  Static    07h00m15s  5
    172.16.1.2          1
172.16.22.0/24         Remote  BGP VPN   00h00m15s  170
    192.0.2.2 (tunneled:RSVP:2)  0
172.16.44.0/24         Remote  BGP VPN   00h09m28s  170
    192.0.2.4 (tunneled)         0
172.16.55.0/24         Remote  BGP VPN   00h09m28s  170
    192.0.2.5 (tunneled)         0
-----
No. of Routes: 8

```

If RSVP is disabled with **configure>router>mpls>shutdown**, the connectivity falls back to LDP as the output shows:

```

*A:PE-1# configure router mpls shutdown
*A:PE-1# show router 1 fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
172.16.1.0/24                              LOCAL
    172.16.1.0 (to-cel)
172.16.2.0/24                              BGP_VPN
    192.0.2.2 (VPRN Label:262142 Transport:LDP)
172.16.4.0/24                              BGP_VPN
    192.0.2.4 (VPRN Label:262142 Transport:LDP)
172.16.5.0/24                              BGP_VPN
    192.0.2.5 (VPRN Label:262142 Transport:LDP)
172.16.11.0/24                             STATIC
    172.16.1.2 (to-cel)
172.16.22.0/24                             BGP_VPN
    192.0.2.2 (VPRN Label:262142 Transport:LDP)
172.16.44.0/24                             BGP_VPN
    192.0.2.4 (VPRN Label:262142 Transport:LDP)
172.16.55.0/24                             BGP_VPN
    192.0.2.5 (VPRN Label:262142 Transport:LDP)
-----
Total Entries : 8

```

If LDP is disabled with “configure router ldp shutdown” the connectivity falls back to GRE as the output shows:

```

*A:PE-1# configure router ldp shutdown
*A:PE-1# show router 1 fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
172.16.1.0/24                              LOCAL
    172.16.1.0 (to-cel)
172.16.2.0/24                              BGP_VPN
    192.0.2.2 (VPRN Label:262142 Transport:GRE)

```

MPLS/GRE Shortcut for BGP NH Resolution within a VRF

```
172.16.4.0/24                               BGP_VPN
  192.0.2.4 (VPRN Label:262142 Transport:GRE)
172.16.5.0/24                               BGP_VPN
  192.0.2.5 (VPRN Label:262142 Transport:GRE)
172.16.11.0/24                              STATIC
  172.16.1.2 (to-cel)
172.16.22.0/24                              BGP_VPN
  192.0.2.2 (VPRN Label:262142 Transport:GRE)
172.16.44.0/24                              BGP_VPN
  192.0.2.4 (VPRN Label:262142 Transport:GRE)
172.16.55.0/24                              BGP_VPN
  192.0.2.5 (VPRN Label:262142 Transport:GRE)
```

Total Entries : 8

Conclusion

IGP shortcuts provide a variety of shortcuts in IP, MPLS and IP-VPN scenarios to customers who want to use new options for building routing topologies. Because IGP shortcuts are enabled on a per router basis, SPF computations are independent and irrelevant to other routers, so there is no need to enable shortcuts globally. This network example shows the configuration of IGP shortcuts together with the associated show outputs which can be used for verification and troubleshooting.

Conclusion