Chapter 9

Problem Establishing a GMPLS LSP

This case study describes a problem with establishing a Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP). Specifically, the configuration of the data channel is incorrect because the configuration includes different interface types at both ends of the tunnel. The principles and solution used in this case study also apply to control channel configuration.

The chapter includes a brief summary of GRE tunnels within the context of GMPLS, an example network scenario, and commands to troubleshoot and resolve the problem. (See Table 12.)

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Table 12: Problem Establishing a GRE Tunnel Checklist

Troubleshooting Tasks	Command or Action
Troubleshooting GMPLS and GRE Tunnel on page 186	
Symptom on page 188	show mpls lsp show rsvp session
Cause on page 188	The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.
Troubleshooting Commands on page 188	show mpls lsp extensive show rsvp session detail show link-management peer show link-management te-link show configuration protocols mpls monitor start *filename* show log *filename*
Solution on page 193	Configure both ends of the data channel with the same switching type. show configuration protocols link-management show mpls lsp show link-management te-link
Conclusion on page 194	Both ends of a GMPLS data must be the same encapsulation or interface type.
Router Configurations on page 195	show configuration | no-more

Troubleshooting GMPLS and GRE Tunnel

The logical control channel for GMPLS must be a point-to-point link and must have some form of IP reachability. On broadcast interfaces or when there are multiple hops between control channel peers, use a GRE tunnel for the control channel. For more detailed information on GMPLS and GRE tunnels see the *JUNOS MPLS Applications Configuration Guide* and the *JUNOS Feature Guide*.

A tunnel PIC is *not* required to configure a GRE tunnel for the GMPLS control channel. Instead, use the software-based gre interface, rather than the hardware-based gr-*fpc/pic/port* interface.

CAUTION: Due to restrictions to the software-based gre interface, the GMPLS control channel is the only supported use of the software-based gre interface. Any other use is expressly unsupported and might cause an application failure.

The following example shows a basic gre interface configuration. In this case, the tunnel source is the loopback address of the local router and the destination address is the loopback destination of the remote router. Traffic that has a next hop of the tunnel destination will use the tunnel. The tunnel is not automatically used by all the traffic passing through the interface. Only traffic with the tunnel destination as the next hop uses the tunnel.

Sample Output

```
user@R1> show configuration interfaces
[...Output truncated...]
gre {
    unit 0 {
        tunnel {
            source 10.0.12.13;
            destination 10.0.12.14;
        }
        family inet {
            address 10.35.1.6/30;
        }
        family mpls;
    }
}
```

Sample Output

The following sample output for the show interfaces command shows the encapsulation type and header, the maximum speed, packets through the logical interface, the destination, and logical address.

```
user@R1> show interfaces gre
Physical interface: gre, Enabled, Physical link is Up
  Interface index: 10, SNMP ifIndex: 8
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
    Input packets : 0
    Output packets: 0

  Logical interface gre.0 (Index 70) (SNMP ifIndex 47)
    Flags: Point-To-Point SNMP-Traps 0x4000
    IP-Header 10.0.12.14:10.0.12.13:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    Input packets : 171734
    Output packets: 194560
    Protocol inet, MTU: 1476
```

```
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.35.1.4/30, Local: 10.35.1.6, Broadcast: 10.35.1.7
Protocol mpls, MTU: 1464
  Flags: None
```

The following are various requirements when you configure a GMPLS LSP using a GRE tunnel:

- The data channel must start and end on the same type of interface.
- The control channel can be a GRE tunnel that starts and ends on the same or different interface type.
- The GRE tunnel must be configured indirectly with the peer-interface *peer-name* statement at the [edit protocol ospf] hierarchy level.
- The GRE interface must be disabled at the [edit protocols ospf] and [edit protocols rsvp] hierarchy levels.
- Data and control channels must be defined correctly in the LMP configuration .
- It is optional to disable Constrained Shortest Path First (CSPF) with the no-cspf statement.

This case focuses on the incorrect configuration of the endpoints of the GRE tunnel. However, you can use a similar process and commands to diagnose other GRE tunnel problems. Figure 18 illustrates a network topology with MPLS tunneled through a GRE interface.

Figure 18: GMPLS Network Topology

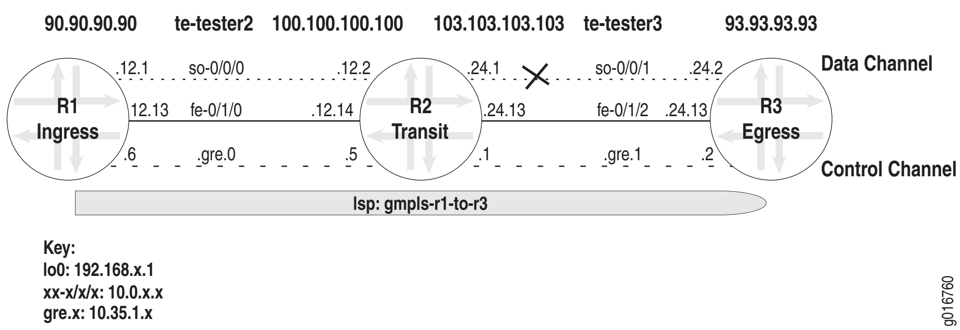

The MPLS network topology in Figure 18 shows Juniper Networks routers configured with a GRE tunnel that consists of the following components:

- A strict GMPLS LSP path from the ingress router to the egress router.
- On the ingress router, CSPF disabled with the no-cspf statement at the [edit protocol mpls label-switched-path *lsp-name*] hierarchy level.
- Traffic-engineering links and control channels within the peer statement at the [edit protocols link-management] hierarchy level on all routers.
- OSPF and OSPF traffic engineering configured on all routers.

- A reference to the peer-interface in both OSPF and RSVP on all routers.
- A switching-type problem between R2 and R3.

Symptom

The LSP in the network shown in Figure 18 on page 187 is down, as indicated by the output from the show mpls lsp and show rsvp session commands, which display very similar information. The show mpls lsp command shows all LSPs configured on the router, as well as all transit and egress LSPs. The show rsvp session command displays summary information about RSVP sessions. You can use either command to verify the state of the LSP. In this case, LSP gmpls-r1-to-r3 is down (Dn).

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.4.1     192.168.1.1     Dn     0 -                      gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R1> show rsvp session
Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.4.1     192.168.1.1     Dn       0  0  -       -        - gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 0, Down 1

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Cause

The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.

Troubleshooting Commands

The JUNOS software includes commands that are useful when troubleshooting a problem. This section provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the problem.

You can use the following commands when troubleshooting a GMPLS problem:

user@host> **show mpls lsp extensive**
user@host> **show rsvp session detail**
user@host> **show link-management peer**
user@host> **show link-management te-link**
user@host> **show configuration protocols mpls**
user@host> **monitor start** *filename*

user@host> **show log** *filename*

Sample Output Use the show mpls lsp extensive command on transit router R1 to display detailed information about all LSPs transiting, terminating, and configured on the router.

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Dn, ActiveRoute: 0, LSPname: gmpls-r1-to-r3
  Bidirectional
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: SDH/SONET, Switching type: PSC-1, GPID: IPv4
  Primary   p1              State: Dn
    SmartOptimizeTimer: 180
    8 Dec 20 18:08:02 192.168.4.1: MPLS label allocation failure[3 times]
    7 Dec 20 18:07:53 Originate Call
    6 Dec 20 18:07:53 Clear Call
    5 Dec 20 18:07:53 Deselected as active
    4 Dec 20 18:06:13 Selected as active path
    3 Dec 20 18:06:13 Record Route:  100.100.100.100 93.93.93.93
    2 Dec 20 18:06:13 Up
    1 Dec 20 18:06:13 Originate Call
  Created: Wed Dec 20 18:06:12 2006
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output for the show mpls lsp extensive command shows an error message (MPLS label allocation failure) in the log section of the output. This LSP event indicates that the MPLS protocol or the family mpls statement were not configured properly. When the LSP event is preceded by an IP address, the address is typically the router that has the MPLS configuration error. In this case, it appears that the router with the lo0 address of 192.168.4.1 (R3) has an MPLS configuration error.

Sample Output Use the show rsvp session detail command to display detailed information about RSVP sessions.

```
user@R1> show rsvp session detail
Ingress RSVP: 1 sessions

192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left:    -, Since: Wed Dec 20 18:07:53 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 0
  PATH sentto: 10.35.1.5 (tester2) 3 pkts
```

```
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> ...incomplete
Total 1 displayed, Up 0, Down 1

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output for the show rsvp session detail command shows that LSP gmpls-r1-to-r3 is down (LSPstate: Dn). The route record is incomplete, indicating a problem with the explicit route 100.100.100.100 93.93.93.93. The address 100.100.100.100 is the data channel on R2 so-0/0/0, and the address 93.93.93.93 is the data channel on R3.

Sample Output Use the show link-management peer command to display MPLS peer link information.

```
user@R1> show link-management peer
Peer name: tester2, System identifier: 48428
 State: Up, Control address: 10.35.1.5
   Control-channel                   State
   gre.0                             Active
  TE links:
   tester2

user@R2> show link-management peer
Peer name: tester2, System identifier: 48428
 State: Up, Control address: 10.35.1.6
   Control-channel                   State
   gre.0                             Active
  TE links:
   te-tester2

Peer name: tester3, System identifier: 48429
 State: Up, Control address: 10.35.1.2
   Control-channel                   State
   gre.1                             Active
  TE links:
   te-tester3

user@R3> show link-management peer
Peer name: tester3, System identifier: 48429
 State: Up, Control address: 10.35.1.1
   Control-channel                   State
   gre.0                             Active
  TE links:
   te-tester3
```

What It Means The sample output from all routers in the example network in Figure 18 on page 187 for the show link-management peer command shows that all control channels are up and active. A detailed analysis of the output shows the following information:

- Name of the peer, tester2 or tester3, which is the same on neighboring routers for ease of troubleshooting.
- Internal identifier for the peer, 48428 for tester2 and 48429 for tester3. The internal identifier is a range of values from 0 through 64,000.

- The state of the peer, which can be up or down. In this case, all peers are up.
- The address to which a control channel is established, for example, 10.35.1.5.
- The state of the control channel, which can be up, down, or active.
- The traffic-engineered links that are managed by their peer, indicating that control channel gre.0 is managed by tester3.

Sample Output Use the show link-management te-link command to display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

```
user@R1> show link-management te-link
 TE link name: tester2, State: Up
  Local identifier: 2005, Remote identifier: 21253, Local address: 90.90.90.90,
Remote address: 100.100.100.100,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/0    Up       21253      21253     155.52Mbps  Yes  gmpls-r1-to-r3

user@R2> show link-management te-link
 TE link name: te-tester2, State: Up
  Local identifier: 7002, Remote identifier: 22292, Local address:
100.100.100.100, Remote address: 90.90.90.90,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/0    Up       21253      21253     155.52Mbps  Yes  gmpls-r1-to-r3
TE link name: te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address:
103.103.103.103, Remote address: 93.93.93.93,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/1    Up       21252      21252     155.52Mbps  Yes  gmpls-r1-to-r3

user@R3> show link-management te-link
 TE link name: te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address: 93.93.93.93,
Remote address: 103.103.103.103,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 0bps, Maximum
bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/1    Dn       21252      21252     155.52Mbps   No
```

What It Means The sample output for the show link-management te-link command issued on the three routers in the network in Figure 18 on page 187 shows the resources allocated to the traffic-engineered links te-tester2 and te-tester3. The resources are the SONET interfaces so-0/0/0 and so-0/0/1. On R1 and R2, the SONET interfaces are used for the LSP gmpls-r1-to-r3, as indicated by Yes in the Used field. However, the SONET interface so-0/0/1 on R3 is down (Dn) and is not used for the LSP (Used No). Further investigation is required to discover why the SONET interface on R3 is down.

Sample Output Use the show log *filename* command to display the contents of the specified log file. In this case, the log file rsvp.log is configured at the [edit protocols rsvp traceoptions] hierarchy level. When the log file is configured, you must issue the monitor start *filename* command to begin logging messages to the file.

```
user@R1> show configuration protocols rsvp
traceoptions {
    file rsvp.log size 3m world-readable;
    flag state detail;
    flag error detail;
    flag packets detail;
}

user@R1> monitor start rsvp.log
```

NOTE: The find Error option entered after the pipe (|) searches the output for an instance of the term *Error*.

```
user@R3> show log rsvp.log | find Error
Dec 28 17:23:32   Error    Len 20 Session preempted flag 0 by 192.168.4.1 TE-link 103.103.103.103
[...Output truncated...]
Dec 28 17:23:32 RSVP new resv state,session 192.168.4.1(port/tunnel ID 46115 Ext-ID 192.168.1.1)Proto 0
Dec 28 17:23:32         RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32         RSVP->LMP request -  resource for LSP gmpls-r1-to-r3
Dec 28 17:23:32         LMP->RSVP resource request gmpls-r1-to-r3 failed cannot find resource  encoding
type SDH/SONET remote label 21252 bandwidth bw[0
Dec 28 17:23:32         RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32 RSVP originate PathErr 192.168.4.1->192.168.2.1 MPLS label allocation failure LSP
gmpls-r1-to-r3(2/46115)
Dec 28 17:23:32 RSVP send PathErr 192.168.4.1->192.168.2.1 Len=196 tester3
Dec 28 17:23:32   Session7 Len 16 192.168.4.1(port/tunnel ID 46115 Ext-ID 192.168.1.1) Proto 0
Dec 28 17:23:32   Hop      Len 20 192.168.4.1/0x086e4770 TE-link 103.103.103.103
Dec 28 17:23:32   Error    Len 20 MPLS label allocation failure flag 0 by 192.168.4.1 TE-link
103.103.103.103
Dec 28 17:23:32   Sender7  Len 12 192.168.1.1(port/lsp ID  2)
Dec 28 17:23:32   Tspec    Len 36 rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
Dec 28 17:23:32   ADspec   Len 48 MTU 1500
Dec 28 17:23:32   RecRoute Len 20  103.103.103.103 90.90.90.90
Dec 28 17:23:32   SuggLabel Len  8  21252
Dec 28 17:23:32   UpstrLabel Len  8  21252
```

What It Means The sample output from the egress router R3 for the show log rsvp.log command is a snippet taken from the log file. The snippet shows a Link Management Protocol (LMP) resource request for the LSP gmpls-r1-to-r3. The request has problems with the encoding type (SDH/SONET), indicating a possible error with the SONET interface connecting R2 and R3. Further investigation of the configuration of the LMP on R2 and R3 is required.

Sample Output Use the show configuration *statement-path* command to display a specific configuration hierarchy; in this instance, link-management.

```
user@R2> show configuration protocols link-management
te-link te-tester2 {
    local-address 100.100.100.100;
    remote-address 90.90.90.90;
    remote-id 22292;
    interface so-0/0/0 {
        local-address 100.100.100.100;
        remote-address 90.90.90.90;
```

```
        remote-id 21253;
    }
}
te-link te-tester3 {
    local-address 103.103.103.103;
    remote-address 93.93.93.93;
    remote-id 21254;
    interface so-0/0/1 {
        local-address 103.103.103.103;
        remote-address 93.93.93.93;
        remote-id 21252;
    }
}
peer tester2 {
    address 10.35.1.6;
    control-channel gre.0;
    te-link te-tester2;
}
peer tester3 {
    address 10.35.1.2;
    control-channel gre.1;
    te-link te-tester3;
}

user@R3> show configuration protocols link-management
te-link te-tester3 {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
    }
    interface at-0/3/1 {
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}
```

What It Means The sample output from transit router R2 and ingress router R3 for the show configuration protocols link-management command shows that the interface type on the two routers is different. The resource allocated to te-tester3 on transit router R2 is a SONET interface, while the resource allocated to te-tester3 on egress router R3 is an ATM interface. The interface type on each end of the data or control channels must be of the same type. In this case, both ends should be SONET or ATM.

Solution

The solution to the problem of different interface or encapsulation types at either end of the GMPLS LSP is to make sure that the interface type is the same at both ends. In this case, the ATM interface was deleted from the link-management configuration on R3, and a SONET interface was configured instead.

The following commands illustrate the correct configuration and commands to verify that the GMPLS LSP is up and using the data channel:

user@R3> **show configuration protocols link-management**
user@R3> **show mpls lsp**
user@R3> **show link-management te-link**

Sample Output

```
user@R3> show configuration protocols link-management
te-link te-tester3 {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
    interface so-0/0/1 { # SONET interface replaces the incorrect ATM interface
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}

user@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.4.1     192.168.1.1     Up       0  1 FF   21252        - gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show link-management te-link
 TE link name: te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address: 93.93.93.93,
Remote address: 103.103.103.103,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/1    Up       21252      21252     155.52Mbps  Yes  gmpls-r1-to-r3
```

What It Means The sample output for the show protocols link-management, show mpls lsp, and show link-management te-link commands from ingress router R3 show that the problem is solved. LMP is correctly configured, and the LSP gmpls-r1-to-r3 is up and using the data channel so-0/0/1.

Conclusion

In conclusion, both ends of a GMPLS data channel must be the same encapsulation or interface type. This case illustrates the correct configuration of the data channel. The principles are the same for the control channel.

Router Configurations

Purpose Output that shows the configurations of the ingress router in the network. The no-more option entered after the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.1/32 {
                    destination 10.0.12.2;
                }
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    gre {
        unit 0 {
            tunnel {
                source 10.0.12.13;
                destination 10.0.12.14;
            }
            family inet {
                address 10.35.1.6/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        /* corporate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
```

```
                no-readvertise;
            }
            /* old lab nets */
            route 192.168.0.0/16 {
                next-hop 192.168.71.254;
                retain;
                no-readvertise;
            }
            route 0.0.0.0/0 {
                discard;
                retain;
                no-readvertise;
            }
        }
        router-id 192.168.1.1;
        autonomous-system 65432;
    }
    protocols {
        rsvp {
            traceoptions {
                file rsvp.log size 3m world-readable;
                flag state detail;
                flag error detail;
                flag packets detail;
            }
            interface fxp0.0 {
                disable;
            }
            interface all;
            interface lo0.0;
            interface gre.0 {
                disable;
            }
            peer-interface tester2;
        }
        mpls {
            label-switched-path gmpls-r1-to-r3 {
                from 192.168.1.1;
                to 192.168.4.1;
                lsp-attributes {
                    switching-type psc-1;
                    encoding-type sonet-sdh;
                }
                no-cspf;
                primary p1;
            }
            path p1 {
                100.100.100.100 strict;
                93.93.93.93 strict;
            }
            interface all;
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface lo0.0;
                interface fe-0/1/0.0;
                interface fxp0.0 {
                    disable;
                }
                interface gre.0 {
                    disable;
                }
```

```
                peer-interface tester2;
            }
        }
        link-management {
            te-link tester2 {
                local-address 90.90.90.90;
                remote-address 100.100.100.100;
                remote-id 21253;
                interface so-0/0/0 {
                    local-address 90.90.90.90;
                    remote-address 100.100.100.100;
                    remote-id 21253;
                }
            }
            peer tester2 {
                address 10.35.1.5;
                control-channel gre.0;
                te-link tester2;
            }
        }
    }
```

Sample Output The following sample output is for transit router R2:

```
user@R2> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.2/32 {
                    destination 10.0.12.1;
                }
            }
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.1/32 {
                    destination 10.0.24.2;
                }
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.14/30;
            }
            family mpls;
        }
    }
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.24.13/30;
            }
            family mpls;
        }
```

```
        }
        fxp0 {
            unit 0 {
                family inet {
                    address 192.168.70.144/21;
                }
            }
        }
        gre {
            unit 0 {
                tunnel {
                    source 10.0.12.14;
                    destination 10.0.12.13;
                }
                family inet {
                    address 10.35.1.5/30;
                }
                family mpls;
            }
            unit 1 {
                tunnel {
                    source 10.0.24.13;
                    destination 10.0.24.14;
                }
                family inet {
                    address 10.35.1.1/30;
                }
                family mpls;
            }
        }
        lo0 {
            unit 0 {
                family inet {
                    address 192.168.2.1/32;
                }
            }
        }
    }
    routing-options {
        static {
            route 172.16.0.0/12 {
                next-hop 192.168.71.254;
                retain;
                no-readvertise;
            }
            route 192.168.0.0/16 {
                next-hop 192.168.71.254;
                retain;
                no-readvertise;
            }
            route 0.0.0.0/0 {
                discard;
                retain;
                no-readvertise;
            }
        }
        router-id 192.168.2.1;
        autonomous-system 65432;
    }
    protocols {
        rsvp {
            traceoptions {
                file rsvp.log size 3m world-readable;
```

```
            flag packets detail;
            flag state detail;
            flag error detail;
        }
        interface fxp0.0;
        interface lo0.0;
        interface all;
        interface gre.0 {
            disable;
        }
        peer-interface tester2;
        peer-interface tester3;
    }
    mpls {
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface fxp0.0 {
                disable;
            }
            interface gre.0 {
                disable;
            }
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface gre.1 {
                disable;
            }
            peer-interface tester2;
            peer-interface tester3;
        }
    }
    link-management {
        te-link te-tester2 {
            local-address 100.100.100.100;
            remote-address 90.90.90.90;
            remote-id 22292;
            interface so-0/0/0 {
                local-address 100.100.100.100;
                remote-address 90.90.90.90;
                remote-id 21253;
            }
        }
        te-link te-tester3 {
            local-address 103.103.103.103;
            remote-address 93.93.93.93;
            remote-id 21254;
            interface so-0/0/1 {
                local-address 103.103.103.103;
                remote-address 93.93.93.93;
                remote-id 21252;
            }
        }
        peer tester2 {
            address 10.35.1.6;
            control-channel gre.0;
            te-link te-tester2;
        }
        peer tester3 {
            address 10.35.1.2;
```

```
                control-channel gre.1;
                te-link te-tester3;
            }
        }
    }
```

Sample Output The following sample output is for egress router R3:

```
user@R3> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/32;
            }
            family mpls;
        }
    }
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.24.14/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    gre {
        unit 0 {
            tunnel {
                source 10.0.24.14;
                destination 10.0.24.13;
            }
            family inet {
                address 10.35.1.2/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
```

```
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.4.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag packets detail;
            flag error;
            flag state;
            flag lmp;
        }
        interface fxp0.0 {
            disable;
        }
        interface all;
        interface lo0.0;
        interface gre.0 {
            disable;
        }
        peer-interface tester3;
    }
    mpls {
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fxp0.0 {
                disable;
            }
            interface fe-0/1/2.0;
            interface gre.0 {
                disable;
            }
            interface lo0.0;
            peer-interface tester3;
        }
    }
    link-management {
        te-link te-tester3 {
            local-address 93.93.93.93;
            remote-address 103.103.103.103;
            remote-id 21254;
            interface so-0/0/1 {
                local-address 93.93.93.93;
                remote-address 103.103.103.103;
                remote-id 21252;
            }
        }
        peer tester3 {
            address 10.35.1.1;
            control-channel gre.0;
```

```
            te-link te-tester3;
        }
    }
}
```