

Chapter 6

GMPLS

Generalized Multiprotocol Label Switching (GMPLS) is the next-generation implementation of Multiprotocol Label Switching (MPLS). GMPLS extends the functionality of MPLS to include a wider range of label-switched path (LSP) options for a variety of network devices.

This document assumes you have a general understanding of MPLS, label switching concepts, and GMPLS Phase 1. For more information about MPLS, see the *JUNOS MPLS Applications Configuration Guide*. For more information about GMPLS Phase 1, see the *JUNOS 5.5 Feature Guide* at: <http://www.juniper.net/techpubs/software/junos/junos55/feature-guide55/feature-guide-55.pdf>.

This feature guide covers these topics:

- Overview on page 286
- System Requirements on page 287
- Terms and Acronyms on page 288
- GMPLS Phase 2 Implementation on page 289
- Configuring GMPLS on page 291
 - Example: GMPLS Configuration on page 300
 - Checking Your Work on page 305
- For More Information on page 311
- Revision History on page 312

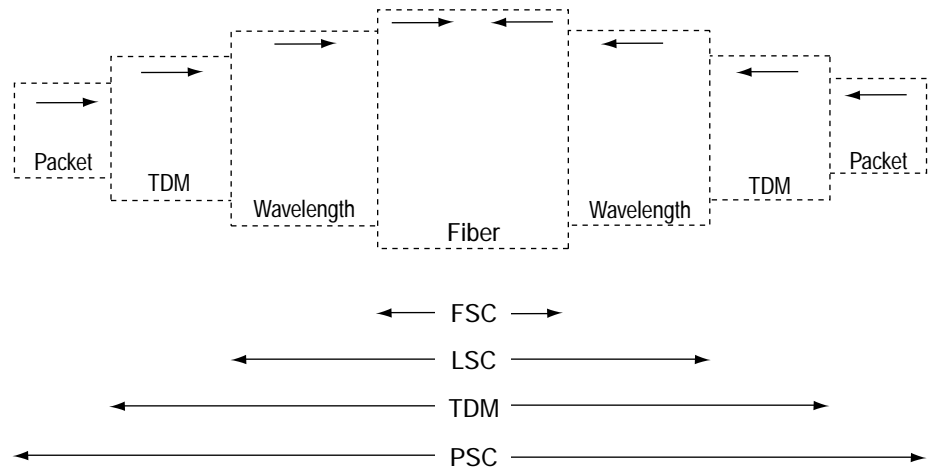
Overview

Traditional MPLS is designed to carry Layer 3 IP traffic by establishing IP-based paths and associating these paths with arbitrarily assigned labels. These labels can either be configured explicitly by a network administrator or dynamically assigned by a protocol such as the Label Distribution Protocol (LDP) or Resource Reservation Protocol (RSVP).

In contrast, GMPLS can carry varying types of Layer 1 through Layer 3 traffic. GMPLS labels and LSPs can be processed at four levels, as depicted in Figure 30. The levels are Fiber-Switched Capable (FSC), Lambda-Switched Capable (LSC), Time-Division Multiplexing Capable (TDM), and Packet-Switched Capable (PSC).

LSPs must start and end on links with the same switching capability. To send an LSP, a label-switched device must communicate with another device at the same layer of the Open System Interconnection (OSI) model. Thus, routers can set up PSC LSPs with other routers at Layer 3, while SONET/SDH add/drop multiplexers (ADMs) can establish TDM LSPs with other ADMs at Layer 1. As seen in Figure 30, a router PSC LSP can be carried over a TDM LSP, a TDM LSP can be carried over a wavelength LSC LSP, and so on.

Figure 30: GMPLS LSP Hierarchy



1760

This extension of the MPLS protocol expands the number of devices that can participate in label-switching. Lower layer devices, such as optical cross-connects (OXC) and SONET/SDH ADMs, can now participate in GMPLS signaling and set up paths to transfer data. Additionally, routers can participate in signaling optical paths across a transport network.

GMPLS labeling is also more flexible than MPLS. A GMPLS label can represent a TDM timeslot, a Dense Wavelength Division Multiplexing (DWDM) wavelength (also known as a lambda), or a physical port number. The labels can be derived from physical components of the network devices, such as interfaces.

There are two service models for GMPLS. Each model determines how much visibility a client node, such as a router, has into the optical core or transport network. The first model is a user-to-network interface (UNI), which is often referred to as the overlay model. The second is known as the peer model. Juniper Networks supports both models.

To enable multilayer LSPs, GMPLS uses the following mechanisms:

- Separation of the control channel from the data channel—A new protocol called Link Management Protocol (LMP) is used to define and manage both control channels and data channels between GMPLS peers. Messages for GMPLS LSP setup are sent from one device to a peer device over an out-of-band control channel. Once the LSP setup is complete and the path is provisioned, the data channel is established and can be used to carry traffic. In GMPLS, the control channel is always separate from the data channel.
- RSVP-TE extensions for GMPLS—RSVP-TE was designed to signal the setup of packet LSPs only. The protocol has been extended to request path setup for non-packet LSPs that use wavelengths, timeslots, and fibers as potential labels.
- OSPF extensions for GMPLS—OSPF was designed to route packets to physical and logical interfaces related to a Physical Interface Card (PIC). This protocol has been extended to route packets to virtual peer interfaces defined in an LMP configuration.
- Bidirectional LSPs—Unlike unidirectional LSP paths found in the standard, packet-based version of MPLS, data can travel both ways between GMPLS devices over a single LSP path. Non-packet LSPs in GMPLS are bidirectional by default.

GMPLS is intended to bridge the gap between the traditional transport infrastructure and the IP layer. Since this protocol is supported by several network industry organizations and standards bodies, GMPLS is designed to enable multivendor interoperability and multilayer functionality. In the near future, routers will be able to make dynamic requests for extra bandwidth on demand from the optical network. Consequently, service providers envision GMPLS as a means to set up optical circuits and services dynamically instead of manually. Many industry professionals are cautiously optimistic regarding the advent of GMPLS, and Juniper Networks is pleased to continue its support for this protocol.

System Requirements

To implement GMPLS Phase 2, your system must meet these minimum requirements:

- JUNOS Release 7.0 or later for graceful teardown of GMPLS LSPs and graceful restart of GMPLS neighbors
- JUNOS Release 5.6 or later for GMPLS Phase 2
- Two Juniper Networks M-series or T-series routing platforms, and two optical cross-connects (OXC) that support GMPLS

Terms and Acronyms

- **Generalized Multiprotocol Label Switching (GMPLS)**—An extension to MPLS that allows data from multiple layers to be switched over label-switched paths (LSPs). GMPLS LSPs are possible between equivalent Layer 1, Layer 2, and Layer 3 devices. For more information about GMPLS and MPLS, see the *JUNOS MPLS Applications Configuration Guide*.
- **control adjacency**—A signaling path between peer devices in a GMPLS network that typically travels across virtual peer interfaces. Protocols are enabled on the control adjacency, which can have one or more associated control channels.
- **control channel**—The actual interfaces where protocol packets are sent and received by GMPLS peers. If more than one control channel is configured, LMP selects which control channel is active.
- **forwarding adjacency**—A forwarding path for sending data between peer devices in a GMPLS network.
- **GMPLS label**—A fiber port, TDM timeslot, DWDM wavelength, or data packet identifier of a GMPLS-enabled device used as a next-hop identifier.
- **GMPLS LSP types**—There are 4 types of LSPs in a GMPLS network:
 - **Fiber-Switched Capable (FSC)**: LSPs are switched between two fiber-based devices, such as optical cross-connects (OXC), that operate at the level of individual fibers.
 - **Lambda-Switched Capable (LSC)**: LSPs are switched between two DWDM devices, such as such as OXC, that operate at the level of individual wavelengths.
 - **TDM-Switched Capable (TDM)**: LSPs are switched between two TDM devices, such as SONET/SDH ADMs.
 - **Packet-Switched Capable (PSC)**: LSPs are switched between two packet-based devices, such as routers or ATM switches.
- **Link Management Protocol**—A GMPLS-related protocol used to define control adjacencies and forwarding adjacencies between peers and to maintain and allocate resources on traffic engineering links (TE links).
- **traffic engineering link (TE link)**—A logical connection between GMPLS-enabled devices. TE links can have addresses or IDs and are associated with certain resources or interfaces. They also have certain inherent attributes, such as encoding-type, switching capability, and bandwidth. Each TE link represents a forwarding adjacency between a pair of devices.

GMPLS Phase 2 Implementation

The major changes between GMPLS Phase 1 and GMPLS Phase 2 are as follows:

- You must configure one or more control channels between peers when you configure LMP (in addition to the existing statements for LMP peers and TE links). The control channels must travel across a point-to-point link or tunnel. To configure, include the `control-channel` statement at the `[edit protocols link-management peer peer-name]` hierarchy level.
- OSPF and RSVP have been extended to allow control adjacencies between peers using virtual peer interfaces. The peer interfaces are derived from LMP and can be used for the control adjacency between peers instead of the physical interfaces. To configure for OSPF, include the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level. To configure for RSVP, include the `peer-interface` statement at the `[edit protocols rsvp]` hierarchy level. However, when you enable peer interfaces, you must disable RSVP and OSPF on all physical control channel interfaces. Alternately, you can omit the physical control channel interfaces when configuring these protocols.
- The Constrained Shortest Path First (CSPF) algorithm has been extended to permit use with non-packet LSPs. In GMPLS Phase 2, the `no-cspf` statement can be omitted from the LSP configuration because it is no longer mandatory. When this statement is omitted, you must configure the signal type attribute for the LSP. For CSPF to work correctly, OSPF extensions for GMPLS need to be implemented on all devices in the GMPLS network.
- LSP paths now can be strict, loose, or dynamic for GMPLS LSPs because TE link information is now exchanged by OSPF. (GMPLS Phase 1 required strict LSP paths.)

The current JUNOS software release supports the following GMPLS functionality:

- Out-of-band signaling controls the setup of LSP paths, enabling a control plane that is separate from the data plane.
- RSVP-TE extensions support additional objects beyond Layer 3 packets, such as ports, timeslots, and wavelengths.
- Link Management Protocol (LMP) creates and maintains a database of TE links, control channels, and peer information. Only the static version of this protocol is supported.
- Bidirectional LSPs are required between non-packet GMPLS devices.
- Several GMPLS label types are defined in RFC 3471, *Generalized MPLS - Signaling Functional Description*. The MPLS, Generalized, SONET/SDH, Suggested, and Upstream label types are supported.
- Generalized labels do not contain a type field because the nodes are expected to know from the context of their connection what type of label to expect. For example, an encoding type, such as Ethernet or SONET/SDH, is determined by the resources in a TE link.

- Traffic parameters facilitate GMPLS bandwidth encoding and SONET/SDH formatting.
- Interface Identification/Errored Interface Identification, UNI-style signaling, and Secondary LSP paths are supported.
- Original channelized interfaces (such as channelized OC12 to DS3, channelized OC3 to T1, and channelized STM1 to E1) support GMPLS signaling.
- GMPLS graceful restart for RSVP LSP paths

The following functionality is *not* supported in this release:

- Notify messages
- RSVP-TE over unnumbered links
- GMPLS routing extensions for IS-IS
- GMPLS link bundling
- Dynamic LMP

For additional information about GMPLS standards, see “For More Information” on page 311.



NOTE: There is not necessarily a one-to-one correspondence between a physical interface and a GMPLS interface. If a GMPLS connection uses a non-channelized physical connector, the GMPLS label can use the physical port ID. However, the label for channelized interfaces often is based on a channel or timeslot. Consequently, it is best not to refer to GMPLS labels as “interfaces.” To avoid confusion, refer to them as TE links and refer to the physical interfaces as resources.

GMPLS Operation

GMPLS requires close interaction between LMP, RSVP, and OSPF. The following sequence of events describes how GMPLS works:

1. LMP notifies RSVP and OSPF of the control peer, the control adjacency, and resources for the TE link.
2. GMPLS extracts the LSP attributes from the configuration and requests RSVP to signal one or more specific paths, specified by the TE link addresses.
3. RSVP determines the local TE link, corresponding control adjacency and active control channel, and transmission parameters (such as IP destination). It requests that LMP allocate a resource from the TE link with the specified attributes. If LMP successfully finds a resource matching the attributes, label allocation succeeds. RSVP sends a **PathMsg** hop-by-hop until it reaches the target router.
4. The target router, on receiving the RSVP **PathMsg**, requests that LMP allocate a resource based on the signaled parameters. If label allocation succeeds, it sends back a **ResvMsg**.
5. If the signaling is successful, an optical path is provisioned.

Configuring GMPLS

You must complete the following tasks to implement GMPLS:

- Configuring Link Management Protocol Traffic Engineering Links on page 292
- Configuring Link Management Protocol Peers on page 293
- Configuring Peer Interfaces in OSPF and RSVP on page 294
- Establishing GMPLS LSP Path Information on page 294
- Defining GMPLS Label-Switched Paths on page 295
- Discovering Local Identifiers and Configuring Remote Identifiers on page 296

The following are optional GMPLS tasks:

- Option: Tearing Down GMPLS LSPs Gracefully on page 298
- Option: Selecting the Peer Model or the Overlay Model for GMPLS on page 298
- Option: GMPLS Graceful Restart on page 299

To view a GMPLS configuration example, see the following sections:

- Example: GMPLS Configuration on page 300
- Checking Your Work on page 305

Configuring Link Management Protocol Traffic Engineering Links

To begin your GMPLS configuration, enable LMP to define the data channel interconnection between devices at the [edit protocols link-management] hierarchy level.

To configure data channels in LMP, include the `te-link te-link-name` statement at the [edit protocols link-management] hierarchy level. Define all TE link options shown. (You will configure `remote-id` statements at the `te-link` and `interface` levels later.) We recommend that you use a different IP address and mask on your TE link addresses from the ones configured on your physical interfaces. This way, you can identify which addresses are physical and which ones belong to the TE link.

```
[edit]
protocols {
  link-management {
    te-link te-link-name {      # Collection of physical ports or timeslots.
      local-address ip-address; # Local IP address associated with the TE link.
      remote-address ip-address; # Remote IP address mapped to the TE link.
      interface interface-name { # Interface used for data transfer.
        local-address ip-address; # Local IP address for the TE link.
        remote-address ip-address; # Remote IP address for the TE link.
      }
    }
  }
}
```

Configuring Link Management Protocol Peers

After you set up TE links, configure LMP network peers with the `peer` statement at the `[edit protocols link-management]` hierarchy level. A peer is the network device that your router communicates with when setting up the control and data channels. Often, the peer is an OXC. Designate a peer name, configure the peer's router ID as the address (often a loopback address), specify the interface that will be used as a control channel, and apply the TE link to be associated with this peer.

You can configure one or more control channels for a peer. The control channels must have point-to-point connectivity with the peer (for example, you can use a point-to-point link or a tunnel). You can also configure the generic routing encapsulation (GRE) tunnel interface of the Routing Engine as a control channel. To configure, include the `control-channel` statement at the `[edit protocols link-management peer peer-name]` hierarchy level. Without a control channel, the configuration will fail to commit.

```
[edit]
protocols {
  link-management {
    peer peer-name { # Configure the name of your network peer.
      address ip-address; # Include the router ID of the peer.
      control-channel interface; # Specify the interface for the control channel.
      te-link te-link-name; # Assign a TE link to this peer.
    }
  }
}
```



NOTE: Although you can configure the `gre-` tunnel interface on a Routing Engine as a control channel, this interface is not supported nor is it configurable for other applications.

Configuring Peer Interfaces in OSPF and RSVP

After you establish LMP peers, add peer interfaces to OSPF and RSVP. A peer interface is a virtual interface used to support a control adjacency between two peers. OSPF and RSVP form adjacencies between peers by using the peer interfaces instead of the physical interfaces.

Because actual protocol packets are sent and received by peer interfaces, the peer interfaces can be signaled and advertised to peers like any other interface enabled for RSVP and OSPF. The peer interface name must match the peer name configured in LMP. To configure RSVP signaling for LMP peers, include the `peer-interface` statement at the `[edit protocols rsvp]` hierarchy level. To configure OSPF routing for LMP peers, include the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level.

```
[edit]
protocols {
  rsvp {
    peer-interface peer-name { # Configure the name of your LMP peer.
    }
  }
  ospf {
    area area-number {
      peer-interface peer-name { # Configure the name of your LMP peer.
      }
    }
  }
}
```



NOTE: When adding the virtual peer interfaces to RSVP and OSPF, do not configure the corresponding physical control channel interface in either protocol. If the `interface all` option is used, you must disable the protocols manually on the control channel interface.

- To disable OSPF, use the `disable` statement at the `[edit protocols ospf area area-number interface interface-name]` hierarchy level.
- To disable RSVP, use the `disable` statement at the `[edit protocols rsvp interface interface-name]` hierarchy level.

Establishing GMPLS LSP Path Information

When you configure LSP paths for GMPLS, you must use the TE link remote address as your next-hop address. When CSPF is supported, you can use any path option you wish. However, when CSPF is disabled with the `no-cspf` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level, you must use strict paths.

```
[edit]
protocols {
  mpls {
    path path-name {
      next-hop-address (strict | loose);
    }
  }
}
```

Defining GMPLS Label-Switched Paths

Next, define LSP attributes at the [edit protocols mpls label-switched-path] hierarchy level. To enable the proper GMPLS switching parameters, configure the attributes appropriate for your network connection. The default values, which are also appropriate for standard MPLS, are `ipv4` for `gp-id`, `none` for `signal-bandwidth`, and `psc-1` for `switching-type`.



NOTE: In JUNOS Release 5.6 or later, the `signal-bandwidth` statement replaces the `signal-type` statement. Also, virtual tributary (VT) 1.5 and 2.0 SONET/SDH bandwidth options are available at the [edit protocols mpls label-switched-path lsp-name lsp-attributes signal-bandwidth] hierarchy level.

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from ip-address;
      to ip-address;
      primary path-name;
      secondary path-name;
      no-cspf;           # This statement to disable CSPF is optional.
      lsp-attributes {  # Attributes determine the selection of an LSP.
        gp-id type;    # Payload type, such as Ethernet or PPP.
        signal-bandwidth type; # Bandwidth encoding, such as DS3 or STM1.
        switching-type type; # Switching method, such as psc-1 or lambda.
      }
    }
  }
}
```



NOTE: Because MPLS and GMPLS use the same configuration hierarchy for LSPs, it is helpful to know which LSP attributes control LSP functionality. Standard MPLS packet-switched LSPs are unidirectional, while GMPLS non-packet LSPs are bidirectional.

If you use the default packet switching type of `psc-1`, your LSP becomes unidirectional. To enable a GMPLS bidirectional LSP, you must select a non-packet switching type option, such as `lambda`, `fiber`, or `ethernet`, at the [edit mpls label-switched-path *lsp-name* lsp-attributes] hierarchy level.

Discovering Local Identifiers and Configuring Remote Identifiers

Once LMP is enabled on a router, the router automatically assigns two local IDs: one at the `te-link` level, the other at the `interface` level. You must configure these port-to-label mappings manually for LMP on both the router and its peer. To configure, set the local IDs of one device (such as the router) as remote IDs on the peer device (such as an OXC) with the `remote-id` statement at the `[edit protocols link-management te-link te-link-name]` and `[edit protocols link-management te-link te-link-name interface interface-name]` hierarchy levels.

You can view the TE link and interface local IDs by using the `show link-management te-link` command. Once you have learned these IDs, configure them as `remote-id` statements at the corresponding `te-link` and `interface` levels on the peer.

Because peers vary, check with your OXC vendor for the configuration statements and location of the local ID information for your specific optical peer device. If you do not manage the peer device, ask the peer's administrator to enable LMP and generate the IDs for you. GMPLS will not work unless these local IDs from both the router and the peer are configured as remote IDs on the opposite device.

To disable an entire TE link for administrative purposes, include the `disable` statement at the `[edit protocols link-management te-link te-link-name]` hierarchy level. To disable an interface within a TE link, include the `disable` statement at the `[edit protocols link-management te-link te-link-name interface interface-name]` hierarchy level.

```
[edit]
protocols {
  link-management {
    te-link te-link-name {
      disable;           # Disable the entire TE link.
      remote-id id-number; # TE link ID number of the peer device.
      interface interface-name { # Name of the interface used for data transfer.
        disable;       # Disable an interface in the TE link.
        remote-id id-number; # ID number of the remote device.
      }
    }
  }
}
```

Figure 31: TE Link and Interface ID Example

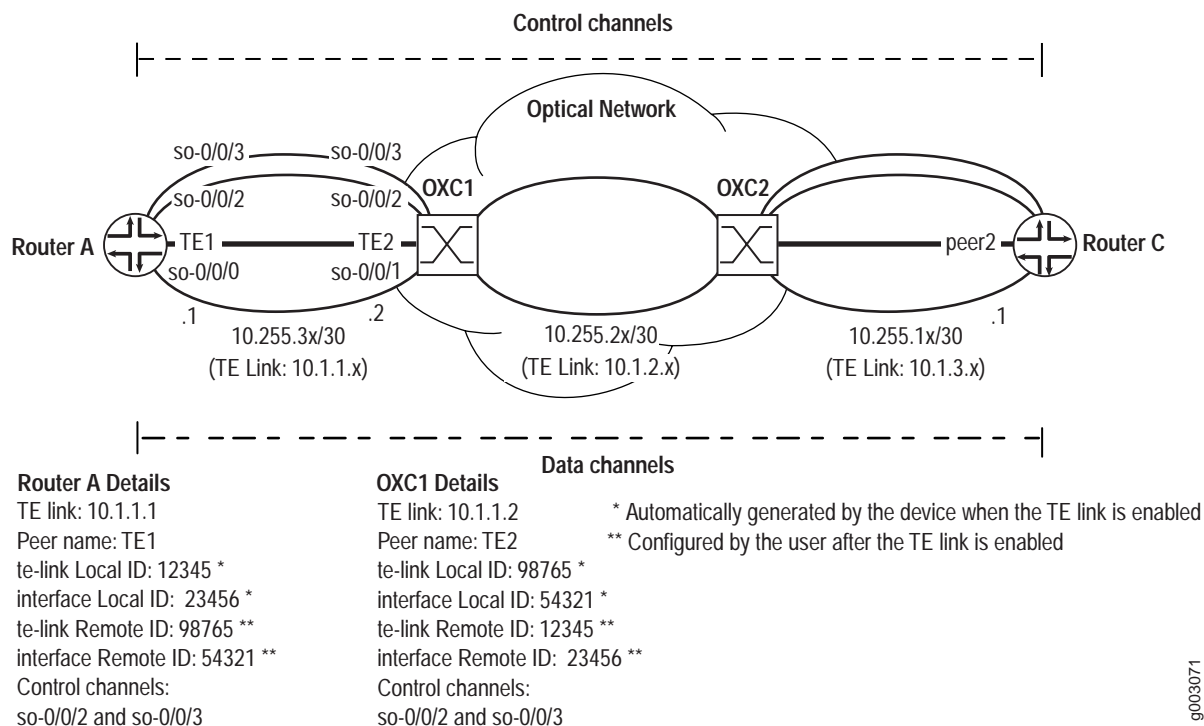


Figure 31 shows where the IDs come from and where you must assign them. This example highlights the connections between Router A and OXC1, but the same configuration concepts apply to all pairs of peers.

First, you configure a TE link named TE1 on Router A, which contains the local address 10.1.1.1, remote address 10.1.1.2, data channel interface so-0/0/0, and control channel interfaces so-0/0/2 and so-0/0/3. You also configure a TE link named TE2 on OXC1, which contains the local address 10.1.1.2, remote address 10.1.1.1, data channel interface so-0/0/1, and control channel interfaces so-0/0/2 and so-0/0/3. When the TE links are enabled on Router A and OXC1, these two peer devices each generate two local IDs: one for the TE link itself and one for the logical interface.

If Router A has a local ID of 12345 for its TE link and a local ID of 23456 for its interface, you must configure 12345 as the TE link remote-ID and 23456 as the interface remote-ID on the TE2 TE link of OXC1. Similarly, if OXC1 has local IDs of 98765 for its TE link and 54321 for its interface, you configure Router A's TE1 TE link with 98765 as the TE link remote-ID and 54321 as the interface remote-ID.

To complete the full data path, you need to enable LMP on each link in the path. This means you must configure remote-ID and local-ID pairs between linked devices. In Figure 31, you would enable LMP three times: between Router A and OXC1, from OXC1 to OXC2, and between OXC2 and Router C.

Option: Tearing Down GMPLS LSPs Gracefully

You can tear down a non-packet GMPLS LSP in a two-step process that gracefully withdraws the RSVP session used by the LSP. For all neighbors that support graceful teardown, a request for the teardown is sent by the routing platform to the destination endpoint for the LSP and all RSVP neighbors in the path. The request is included within the `ADMIN_STATUS` field of the RSVP packet. When neighbors receive the request, they prepare for the RSVP session to be withdrawn. A second message is sent by the routing platform to complete the teardown of the RSVP session. If a neighbor does not support graceful teardown, the request is handled as a standard session teardown rather than a graceful one.

To perform a graceful teardown of a GMPLS LSP RSVP session, issue the `clear rsvp session gracefully` command. Optionally, you can specify the source and destination address of the RSVP session, the LSP identifier of the RSVP sender, and the tunnel identifier of the RSVP session. To use these qualifiers, include the `connection-source`, `connection-destination`, `lsp-id`, and `tunnel-id` options when you issue the `clear rsvp session gracefully` command.

You can also configure the amount of time that the routing platform waits for neighbors to receive the graceful teardown request before initiating the actual teardown. To configure, include the `graceful-deletion-timeout` statement at the `[edit protocols rsvp]` hierarchy level. The default graceful deletion timeout value is 30 seconds, with a minimum value of 1 second and a maximum value of 300 seconds. To view the current value configured for graceful deletion timeout, issue the `show rsvp version operational` command.

Option: Selecting the Peer Model or the Overlay Model for GMPLS

Juniper Networks supports both the peer model and the overlay model for GMPLS. To implement the peer model, perform the following configuration tasks:

- Use the default CSPF calculation that is inherent in MPLS LSPs.
- Include a TE link and a control channel within the `peer` statement at the `[edit protocols link-management]` hierarchy level.
- Reference the `peer-interface` in both OSPF and RSVP.
- Configure OSPF and OSPF traffic engineering to connect with the OXC.
- Optionally, define a GMPLS LSP path.

To implement the overlay model, perform the following configuration tasks:

- Disable CSPF by including the `no-cspf` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level.
- Define a strict GMPLS LSP data path to the remote OXC peer.
- Include a TE link and a control channel within the `peer` statement at the `[edit protocols link-management]` hierarchy level.
- Reference the `peer-interface` in RSVP.

Option: GMPLS Graceful Restart

GMPLS supports graceful restart, a mechanism that allows a restarting router to continue forwarding packets to neighbors without interruption. The restarting router relies on its forwarding table temporarily while it receives updates from “helper” neighbors that assist the restarting router in rebuilding its routing table.

To enable graceful restart for all routing protocols including GMPLS, include the `graceful-restart` statement at the `[edit routing-options]` hierarchy level. To disable graceful restart for GMPLS and RSVP, include the `disable` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level. To disable GMPLS and RSVP graceful restart helper capability, include the `helper-disable` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level.

To configure the maximum amount of time the routing platform retains information for restarting neighbors, include the `maximum-helper-recovery-time` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level. The default helper recovery time is 180 seconds, with a minimum value of 1 second and a maximum value of 3600 seconds. To view the current value configured for the helper recovery time, issue the `show rsvp version operational mode` command.

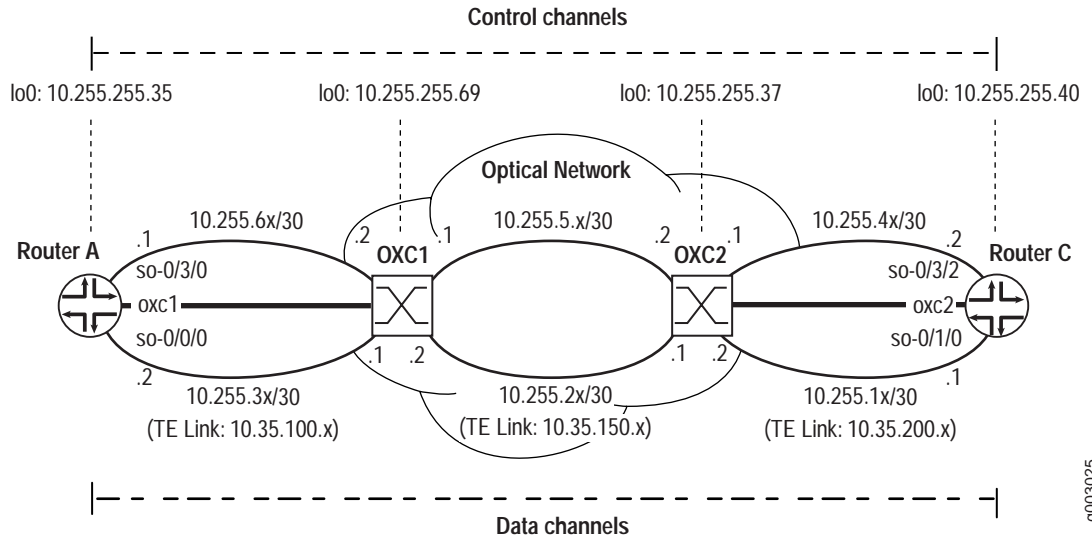
To configure the maximum amount of time the routing platform waits until a neighbor is declared dead, include the `maximum-helper-restart-time` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level. The default helper restart time is 20 seconds, with a minimum value of 1 second and a maximum value of 1800 seconds. To view the current value configured for the helper restart time, issue the `show rsvp version operational` command.

```
[edit]
protocols {
  rsvp {
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time seconds;
      maximum-helper-restart-time seconds;
    }
  }
}
```

For more information about graceful restart, see “Configuring Graceful Restart Options for RSVP, CCC, and TCC” on page 465.

Example: GMPLS Configuration

Figure 32: GMPLS Topology Diagram



In Figure 32, a control channel is established between Router A and OXC1, OXC1 and OXC2, and OXC2 and Router C. A data channel is enabled on a second connection between each pair of devices. The optical network cloud can contain OXCs, ADMs, or other lower-layer devices. In this example, OXC1 and OXC2 are in the direct data path between Routers A and C and the two OXCs have point-to-point connectivity with each other and the directly connected peer routers.

Starting with Router A, configure LMP TE links and peers to create a data channel and a control channel to connect with OXC1. To differentiate the logical TE link from the physical network, the local and remote addresses in the TE link are not related to the IP addresses assigned to the physical interfaces.

When you enable LMP peering on both Router A and OXC1, include the control channel interface as one of the peer statements. Use the name of the peer (in this case, `oxc1`) as the peer interface name when you add the `peer-interface` statement to RSVP at the `[edit protocols RSVP]` hierarchy level and OSPF at the `[edit protocols ospf area area-number]` hierarchy level.

The `peer-interface` statement adds the remote address and local address from your LMP configuration into the routing and signaling processes activated between Router A and OXC1. Make sure the physical control channel is a point-to-point link and has some form of IP reachability through static routes, an interior gateway protocol (IGP), or BGP (this example uses OSPF). Another way to achieve point-to-point links, especially if there are multiple hops between peers, is to use a generic routing encapsulation (GRE) tunnel for the control channel.

Next, configure an MPLS LSP on Router A to reach Router C. For this example, assume your data plane connection uses STM1 and Point-to-Point Protocol (PPP) over a fiber-switched network. Configure these LSP attributes in the LSP. Because this LSP does not use packet switching, a bidirectional LSP is enabled by default. As a result, you do not need to configure a return path LSP on Router C.

Finally, remember to discover the local IDs and configure them on OXC1 with the `remote-id` statement at the `[edit protocols link-management te-link te-link-name]` and `[edit protocols link-management te-link te-link-name interface]` hierarchy levels. For Router A, use the command `show link-management te-link` to find Router A's two local IDs (`te-link` and `interface`); then configure these IDs as remote IDs on OXC1 at the equivalent hierarchy levels.

```

Router A [edit]
interfaces {
  so-0/0/0 {
    description "Data channel to OXC1";
    encapsulation ppp;
    unit 0 {
      family inet {
        address 10.255.3.2/30 {
          destination 10.255.3.1;
        }
      }
      family mpls;
    }
  }
  so-0/3/0 {
    description "Control channel to OXC1";
    encapsulation ppp;
    unit 0 {
      family inet {
        address 10.255.6.1/30 {
          destination 10.255.6.2;
        }
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.35/32;
      }
    }
  }
}

```

```

protocols
  rsvp {
    interface all;
    interface so-0/3/0.0 {
      disable;
    }
    peer-interface oxc1;
  }
  mpls {
    label-switched-path gmpls-lsp1 {
      to 10.255.255.40;
      lsp-attributes {
        signal-bandwidth stm-1;
        switching-type fiber;
        gpid ppp;
      }
      primary path-lsp1;
    }
    path path-lsp1 {
      10.35.100.1 strict; # This example does not disable CSPF,
      10.35.150.1 strict; # so this step is optional.
      10.35.200.1 strict;
    }
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0;
      interface fxp0.0 {
        disable;
      }
    }
    peer-interface oxc1;
  }
}
link-management {
  te-link te-oxc1 {
    local-address 10.35.100.2;
    remote-address 10.35.100.1;
    remote-id 8256;
    interface t3-3/3/0:0 {
      local-address 10.35.100.2;
      remote-address 10.35.100.1;
      remote-id 65536;
    }
  }
  peer oxc1 {
    address 10.255.255.69;
    control-channel so-0/3/0.0;
    te-link te-oxc1;
  }
}
}

```

On OXC1, complete your configuration of the control channel and the TE link data channel to Router A. Refer to your OXC vendor's instructions to configure a TE link on your specific device. Enable LMP peering, configure Router A's local IDs as remote IDs on OXC1, and discover OXC1's local IDs. Finally, configure OXC1's local IDs as remote IDs on Router A.

In the optical network between your OXCs, configure a TE link and a control channel between OXC1 and OXC2. Refer to the OXC vendor's instructions to configure this link. For the example shown in Figure 32 on page 300, you can assume a TE link with an address space of **10.255.150.x/30** has been enabled over a physical network with IP addresses **10.255.2.x/30**. Also, a control channel has been created over the **10.255.4.x/30** link.

On OXC2, configure a TE link to Router A. Refer to your OXC vendor's instructions to configure this TE link on your device. Enable LMP peering, configure Router C's local IDs as remote-IDs on OXC2, and discover OXC2's local IDs. Finally, configure OXC2's local IDs as remote IDs on Router C.

Now you are ready to complete this GMPLS example. On Router C, set up your TE link, LMP peer, and control channel statements to connect to OXC2. As with Router A, the local and remote addresses in the TE link on Router C are not related to the IP addresses assigned to the physical interface.

Next, configure RSVP, MPLS, and OSPF to match the control channel protocols you configured on Router A. You do not need to set up an LSP on Router C because Router A's non-packet LSP is bidirectional by default. Also, because RSVP is enabled for all interfaces and you are using a peer interface, you must disable RSVP on the physical control channel interface **so-0/3/2**.

After you enable LMP on both Router C and OXC2, discover the local IDs and configure them as remote IDs on OXC2. For Router C, use the command **show link-management te-link** to discover Router C's two local IDs (**te-link** and **interface**); then configure these IDs as remote IDs on OXC2 at the equivalent hierarchy levels.

```

Router C [edit]
interfaces {
  so-0/3/2 {
    description "Control channel to OXC2";
    unit 0 {
      family inet {
        address 10.255.4.2/30 {
          destination 10.255.4.1;
        }
      }
      family mpls;
    }
  }
  so-0/1/0 {
    description "Data channel to OXC2";
    encapsulation ppp;
    unit 0 {
      family inet {
        address 10.255.1.1/30 {
          destination 10.255.1.2;
        }
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.255.40/32;
      }
    }
  }
}
protocols
  rsvp {
    interface all;
    interface so-0/3/2.0 {
      disabled;
    }
    peer-interface oxc2;
  }
  mpls {
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface fxp0.0 {
        disable;
      }
      interface lo0.0;
      peer-interface oxc2;
    }
  }
}

```

```

link-management {
  te-link te-oxc2 {
    local-address 10.35.200.1;
    remote-address 10.35.200.2;
    remote-id 41060;
    interface so-0/1/0 {
      local-address 10.35.200.1;
      remote-address 10.35.200.2;
      remote-id 22278;
    }
  }
  peer oxc2 {
    address 10.255.255.37;
    control-channel so-0/3/2.0;
    te-link te-oxc2;
  }
}

```

Checking Your Work

To verify proper operation of GMPLS, you can use the following commands:

- show link-management (te-link | peer)
- show link-management routing (te-link | peer)
- show mpls lsp (bidirectional | unidirectional)
- show mpls lsp (detail | extensive)
- show ospf interface
- show ospf neighbor
- show rsvp interface link-management
- show rsvp session (bidirectional | unidirectional)
- show rsvp session te-link
- show rsvp session detail
- show rsvp neighbor detail
- show ted database extensive
- RSVP traceoptions—lmp flag

The following sections show the output of these commands used with the configuration example:

- Router A Status on page 306
- Router C Status on page 311

Router A Status

After you enter the local-address, remote-address, and interface parameters in TE link `te-oxc1` and commit the changes, the router automatically creates a local ID at the `te-link` and `interface` levels of the `[edit protocols link-management]` hierarchy. To view these IDs, issue the `show link-management te-link` command.

```
user@RouterA> show link-management te-link
TE link name: te-oxc1, State: Up
Local identifier: 8255, Remote identifier: 0, Local address: 10.35.100.2,
Remote address: 10.35.100.1, Encoding: SDH/SONET,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps, Available bandwidth: 0bps
  Name          Local ID Remote ID   Bandwidth In use  LSP
  so-0/0/0      65535     0       155.52Mbps No
```

Once you find these values on Router A, configure them as remote IDs at the same hierarchy levels on OXC1. In this example, 8255 is Router A's local TE link ID (configure this as the TE link remote-ID on OXC1) and 65535 is Router A's local interface ID (configure this as the interface remote-ID on OXC1).

After you configure both remote IDs on both peers, the GMPLS TE links should work. Using the same command as before, you can verify whether the link is functional, with both remote and local IDs in place:

```
user@RouterA> show link-management te-link
TE link name: te-oxc1, State: Up
Local identifier: 8255, Remote identifier: 8256, Local address: 10.35.100.2,
Remote address: 10.35.100.1, Encoding: SDH/SONET,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps, Available bandwidth: 0bps
  Name          Local ID Remote ID   Bandwidth In use  LSP
  so-0/0/0      65535     65536    155.52Mbps Yes      gmp1s-1sp1
```

To further verify proper operation, use the following commands:

```
user@RouterA> show link-management routing peer
Peer name: oxc1, System identifier: 13892
State: Up, Control address: 10.255.255.69
Control-channel          State
so-0/3/0.0              Active

user@RouterA> show link-management routing te-link
TE link name: te-oxc1, State: Up
Local identifier: 8255, Remote identifier: 8256, Local address: 10.35.100.2,
Remote address: 10.35.100.1, Encoding: SDH/SONET,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps, Available bandwidth: 0bps

user@RouterA> show link-management peer
Peer name: oxc1, System identifier: 13892
State: Up, Control address: 10.255.255.69
Control-channel          State
so-0/3/0.0              Active
TE links:
te-oxc1
```

```

user@RouterA> show mpls lsp bidirectional
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P   LSPname
10.255.255.40 10.255.255.35 Up    0 path-lsp1      *   gmpls-lsp1 Bidir
Total 1 displayed, Up 1, Down 0

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

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

user@RouterA> show mpls lsp bidirectional extensive
Ingress LSP: 1 sessions

10.255.255.40
  From: 10.255.255.35, State: Up, ActiveRoute: 0, LSPname: gmpls-lsp1
  Bidirectional
  ActivePath: path-lsp1 (primary)
  LoadBalance: Random
  Signal type: STM-1
  Encoding type: SDH/SONET, Switching type: Fiber, GPID: PPP
  *Primary path-lsp1          State: Up
  Bandwidth: 155.52Mbps
  Computed ERO ($ [L] denotes strict [loose] hops): (CSPF metric: 2)
    10.35.100.1 S 10.35.150.1 S 10.35.200.1 S
  Received RRO:
    10.35.100.1 10.35.150.1 10.35.200.1
  7 Nov 7 15:47:11 Selected as active path
  6 Nov 7 15:47:11 Record Route: 10.35.100.1 10.35.150.1 10.35.200.1
  5 Nov 7 15:47:11 Up
  4 Nov 7 15:47:11 Update LSP Encoding Type
  3 Nov 7 15:47:11 Originate Call
  2 Nov 7 15:47:11 CSPF: computation result accepted
  1 Nov 7 15:46:41 CSPF failed: no route toward 10.255.255.40
  Created: Thu Nov 7 15:46:38 2002
Total 1 displayed, Up 1, Down 0

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

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

```

If you configure an LMP peer interface in OSPF, you can see that this virtual interface is treated as a point-to-point link. To view this, use the `show ospf interface` command.

```

user@RouterA> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR     0.0.0.0   10.255.255.35 0.0.0.0     0
oxc1           PtToPt 0.0.0.0   0.0.0.0     0.0.0.0     1

```

The next command is useful because it indicates whether RSVP is disabled on the control channel. It also shows the state of the reservations on the TE links.

```

user@RouterA> show rsvp interface link-management
RSVP interface: 1 active
oxc1 State Up
Active control channel: so-0/3/0.0 RSVP disabled

    TElink: te-oxc1, Local identifier: 8255
    ActiveResv 1, PreemptionCnt 0
    StaticBW: 155.52Mbps, ReservedBW: 155.52Mbps, AvailableBW: 0bps

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

10.255.255.40
  From: 10.255.255.35, LSPstate: Up, ActiveRoute: 0
  LSPname: gmpls-lsp1, LSPpath: Primary
  Bidirectional, Upstream label in: 27676, Upstream label out: -
  Suggested label received: -, Suggested label sent: 27676
  Recovery label received: -, Recovery label sent: 60444
  Resv style: 1 FF, Label in: -, Label out: 60444
  Time left: -, Since: Thu Nov 7 15:47:11 2002
  Tspec: rate 0bps size 0bps peak 1.544Mbps m 20 M 1500
  Port number: sender 1 receiver 17 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 10.255.255.40 (oxc1) 157 pkts
  RESV rcvfrom: 10.255.255.40 (oxc1) 71 pkts
  Explct route: 10.35.100.1 10.35.150.1 10.35.200.1
  Record route: <self> 10.35.100.1 10.35.150.1 10.35.200.1
Total 1 displayed, Up 1, Down 0

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

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

user@RouterA> show rsvp session bidirectional
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.255.40 10.255.255.35 Up    0 1 FF      -    60444 gmpls-lsp1 Bidir
Total 1 displayed, Up 1, Down 0

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

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

user@RouterA> show rsvp session te-link te-oxc1
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.255.40 10.255.255.35 Up    0 1 FF      -    60444 gmpls-lsp1 Bidir
Total 1 displayed, Up 1, Down 0

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

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

```

```

user@RouterA> show ted database extensive
TED database: 0 ISIS nodes 4 INET nodes
NodeID: 10.255.255.35
  Type: Rtr, Age: 2178 secs, LinkIn: 4, LinkOut: 5
  Protocol: OSPF(0.0.0.0)
    To: 10.255.255.69, Local: 10.35.100.2, Remote: 10.35.100.1
    Metric: 1
    Static BW: 155.52Mbps
    Reservable BW: 155.52Mbps
    Available BW [priority] bps:
      [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
      [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
    Interface Switching Capability Descriptor(1):
      Switching type: Fiber
      Encoding type: SDH/SONET
      Maximum LSP BW [priority] bps:
        [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
        [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
      Minimum LSP BW: 155.52Mbps
      Interface MTU: 2595
NodeID: 10.255.255.37
  Type: Rtr, Age: 2852 secs, LinkIn: 5, LinkOut: 5
  Protocol: OSPF(0.0.0.0)
    To: 10.255.255.69, Local: 10.35.150.1, Remote: 10.35.150.2
    Metric: 1
    Static BW: 622.08Mbps
    Reservable BW: 622.08Mbps
    Available BW [priority] bps:
      [0] 622.08Mbps [1] 622.08Mbps [2] 622.08Mbps [3] 622.08Mbps
      [4] 622.08Mbps [5] 622.08Mbps [6] 622.08Mbps [7] 622.08Mbps
    Interface Switching Capability Descriptor(1):
      Switching type: Fiber
      Encoding type: SDH/SONET
      Maximum LSP BW [priority] bps:
        [0] 622.08Mbps [1] 622.08Mbps [2] 622.08Mbps [3] 622.08Mbps
        [4] 622.08Mbps [5] 622.08Mbps [6] 622.08Mbps [7] 622.08Mbps
      Minimum LSP BW: 622.08Mbps
      Interface MTU: 2597
    To: 10.255.255.40, Local: 10.35.200.2, Remote: 10.35.200.1
    Metric: 1
    Static BW: 155.52Mbps
    Reservable BW: 155.52Mbps
    Available BW [priority] bps:
      [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
      [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
    Interface Switching Capability Descriptor(1):
      Switching type: Fiber
      Encoding type: SDH/SONET
      Maximum LSP BW [priority] bps:
        [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
        [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
      Minimum LSP BW: 155.52Mbps
      Interface MTU: 2600

```

```

NodeID: 10.255.255.40
Type: Rtr, Age: 2854 secs, LinkIn: 2, LinkOut: 2
Protocol: OSPF(0.0.0.0)
  To: 10.255.255.37, Local: 10.35.200.1, Remote: 10.35.200.2
  Metric: 1
  Static BW: 155.52Mbps
  Reservable BW: 155.52Mbps
  Available BW [priority] bps:
    [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
    [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  Interface Switching Capability Descriptor(1):
    Switching type: Fiber
    Encoding type: SDH/SONET
    Maximum LSP BW [priority] bps:
      [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
      [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
    Minimum LSP BW: 155.52Mbps
    Interface MTU: 2600
NodeID: 10.255.255.69
Type: Rtr, Age: 2832 secs, LinkIn: 8, LinkOut: 7
Protocol: OSPF(0.0.0.0)
  To: 10.255.255.35, Local: 10.35.100.1, Remote: 10.35.100.2
  Metric: 1
  Static BW: 155.52Mbps
  Reservable BW: 155.52Mbps
  Available BW [priority] bps:
    [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
    [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  Interface Switching Capability Descriptor(1):
    Switching type: Fiber
    Encoding type: SDH/SONET
    Maximum LSP BW [priority] bps:
      [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
      [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
    Minimum LSP BW: 155.52Mbps
    Interface MTU: 2595
  To: 10.255.255.37, Local: 10.35.150.2, Remote: 10.35.150.1
  Metric: 1
  Static BW: 622.08Mbps
  Reservable BW: 622.08Mbps
  Available BW [priority] bps:
    [0] 622.08Mbps [1] 622.08Mbps [2] 622.08Mbps [3] 622.08Mbps
    [4] 622.08Mbps [5] 622.08Mbps [6] 622.08Mbps [7] 622.08Mbps
  Interface Switching Capability Descriptor(1):
    Switching type: Fiber
    Encoding type: SDH/SONET
    Maximum LSP BW [priority] bps:
      [0] 622.08Mbps [1] 622.08Mbps [2] 622.08Mbps [3] 622.08Mbps
      [4] 622.08Mbps [5] 622.08Mbps [6] 622.08Mbps [7] 622.08Mbps
    Minimum LSP BW: 622.08Mbps
    Interface MTU: 2597

```

```

user@RouterA> show rsvp neighbor detail
RSVP neighbor: 1 learned
Address: 10.255.255.40 via: oxc1 status: Up
Last changed time: 50:52, Idle: 0 sec, Up cnt: 1, Down cnt: 0
Message received: 145
Hello: sent 338, received: 338, interval: 9 sec
Remote instance: 0x643087c7, Local instance: 0x3271e0a4
Refresh reduction: not operational
Link protection: disabled
  Bypass LSP: does not exist, Backup routes: 0, Backup LSPs: 0

```

Router C Status

After you enter the local-address, remote-address, and interface parameters in TE link `te-oxc2` and commit the changes, the router automatically creates a local ID at the `te-link` and `interface` levels of the `[edit protocols link-management]` hierarchy. To view these IDs, issue the `show link-management te-link` command.

```
user@RouterC> show link-management te-link
TE link name: te-oxc2, State: Up
Local identifier: 41059, Remote identifier: 0, Local address: 10.35.200.1,
Remote address: 10.35.200.2, Encoding: SDH/SONET,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps, Available bandwidth: 0bps
Name          Local ID Remote ID   Bandwidth In use  LSP
so-0/1/0      22277    0           155.52Mbps No
```

Once you see what these values are, configure them as remote IDs at the same hierarchy levels on OXC2 where you found them on Router C. In this example, 41059 is Router C's local TE link ID (configure this as the TE link `remote-ID` on OXC2) and 22277 is Router C's local interface ID (configure this as the interface `remote-ID` on OXC2).

After you configure both remote IDs on both peers, the GMPLS TE links should work. Using the same command as before, you can determine whether the link is functional, with both remote and local IDs in place:

```
user@RouterC> show link-management te-link
TE link name: te-oxc2, State: Up
Local identifier: 41059, Remote identifier: 41060, Local address: 10.35.200.1,
Remote address: 10.35.200.2, Encoding: SDH/SONET,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps, Available bandwidth: 0bps
Name          Local ID Remote ID   Bandwidth In use  LSP
so-0/1/0      22277    22278       155.52Mbps Yes    gmp1s-1sp1
```

The other `show` commands operate like those in “Router A Status” on page 306.

For More Information

For additional information about implementing GMPLS, see the following:

- *JUNOS MPLS Applications Configuration Guide*
- RFC 2205, *Resource ReSerVation Protocol (RSVP)*
- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS)—Signaling Functional Description*
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*
- Internet draft draft-ietf-ccamp-gmpls-rsvp-te-ason-02.txt, *Generalized MPLS (GMPLS) RSVP-TE Signalling in support of Automatically Switched Optical Network (ASON)* (expires January 2005)
- Internet draft draft-ietf-ccamp-gmpls-sonet-sdh-08.txt, *GMPLS Extensions for SONET and SDH Control* (expires August 2003)

- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-12.txt, *OSPF Extensions in Support of Generalized MPLS* (expires April 2004)
- Internet draft draft-ietf-mpls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering* (expires January 2003)
- Internet draft draft-ietf-ccamp-imp-10.txt, *Link Management Protocol (LMP)* (expires April 2004)
- Internet draft draft-ietf-mpls-lsp-hierarchy-08.txt, *LSP Hierarchy with Generalized MPLS TE* (expires March 2003)
- Internet draft draft-ietf-ccamp-gmpls-routing-09.txt, *Routing Extensions in Support of Generalized MPLS* (expires April 2004)

Revision History

- 13 June 2005—7.3R1 Release. Richard Hendricks.
- 5 April 2005—Added the steps necessary to configure GMPLS with the peer model and the overlay model, 7.2R1 Release. Richard Hendricks.
- 2 February 2005—7.1R1 Release. Richard Hendricks.
- 6 October 2004—Added support for graceful teardown of GMPLS LSPs and graceful restart of GMPLS neighbors, 7.0R1 Release. Richard Hendricks.
- 6 July 2004—6.4R1 Release. Richard Hendricks.
- 5 April 2004—6.3R1 Release. Richard Hendricks.
- 22 December 2003—6.2R1 Release. Richard Hendricks.
- 22 September 2003—6.1R1 Release. Richard Hendricks.
- 30 June 2003—6.0R1 Release. Richard Hendricks.
- 2 April 2003—5.7R1 Release. Richard Hendricks.
- 27 December 2002—Added OSPF and CSPF updates, 5.6R1 Release. Richard Hendricks.
- 30 September 2002—Added one note for the 5.5R1 Release. Richard Hendricks.
- 19 July 2002—5.4R1 Release. Richard Hendricks.
- 28 June 2002—Revised commands and statements. Richard Hendricks.
- 6 May 2002—Initial document written. Richard Hendricks.