

Chapter 20

GMPLS Configuration Guidelines

To configure GMPLS, you must complete the following tasks:

- Configuring LMP on page 460
- Configuring MPLS LSPs for GMPLS on page 471
- Gracefully Tearing Down GMPLS LSPs on page 473



NOTE: Although you can configure the GMPLS-related statements at the [edit logical-routers *logical-router-name*] hierarchy level, GMPLS is not supported on logical routers.

Configuring LMP

You need to configure the Link Management Protocol (LMP) to define the data channel connection and the control channel connection between devices. Include the following statements at the [edit protocols link-management] hierarchy level:

```
[edit protocols link-management]
peer peer-name {
  address address;
  control-channel control-channel-name;
  lmp-control-channel control-channel-interface {
    remote-address ip-address;
  }
  lmp-protocol {
    hello-interval milliseconds;
    hello-dead-interval milliseconds;
    retransmission-interval milliseconds;
    retry-limit number;
    passive;
  }
  te-link te-link-name;
}
te-link te-link-name {
  disable;
  interface interface-name {
    disable;
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
  }
  label-switched-path label-switched-path-name;
  local-address ip-address;
  remote-address ip-address;
  remote-id id-number;
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size>
    <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

The sections that follow describe how to configure LMP:

- Configuring LMP Traffic Engineering Links on page 461
- Configuring LMP Peers on page 463
- Configuring Peer Interfaces in RSVP and OSPF on page 467
- Configuring MPLS Paths for GMPLS on page 469
- Tracing LMP Traffic on page 470

Configuring LMP Traffic Engineering Links

An LMP traffic engineering link acts as a data channel connection between GMPLS devices.

To configure a traffic engineering link, include the `te-link` statement at the `[edit protocols link-management]` hierarchy level:

```
[edit protocols link-management]
te-link te-link-name {
  disable;
  interface interface-name {
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
  }
  label-switched-path label-switched-path-name;
  local-address ip-address;
  remote-address ip-address;
  remote-id id-number;
}
```

Complete the procedures in the following sections to configure an LMP traffic engineering link:

- Configuring the Local IP Address for the Traffic Engineering Link on page 461
- Configuring the Remote IP Address for the Traffic Engineering Link on page 462
- Configuring the Remote ID for the Traffic Engineering Link on page 462

When you configure a traffic engineering link that contains interfaces for an LMP peer, you must also configure a control channel. However, no control channel is required for a traffic engineering link that contains an LSP. For information on configuring control channels, see “Configuring LMP Peers” on page 463.

Configuring the Local IP Address for the Traffic Engineering Link

Use the `local-address` statement to configure the local IP address associated with the traffic engineering link.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This configuration enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the local IP address for the traffic engineering link, include the `local-address` statement:

```
te-link te-link-name {
  interface interface-name {
    local-address ip-address;
  }
  local-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote IP Address for the Traffic Engineering Link

You need to specify the address of the remote end of the data channel for each traffic engineering link. Use the `remote-address` statement to configure the remote IP address.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the remote IP address for the traffic engineering link, include the `remote-address` statement:

```
te-link te-link-name {
  interface interface-name {
    remote-address ip-address;
  }
  remote-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote ID for the Traffic Engineering Link

The local ID for the traffic engineering link is automatically assigned by LMP. The port identifier and labels for the interfaces (resources) in the traffic engineering link are also assigned automatically. However, you need to explicitly configure the remote ID for the traffic engineering link and the remote ID traffic engineering link interface. The remote ID for the interface must be based on the post-ID assignment of the peer node. The remote IDs are needed for static mapping of remote labels to local labels.

Before you can obtain the remote IDs for the traffic engineering link and traffic engineering link interface on the peer node, you must first configure the LMP peer, as described in “Configuring LMP Peers” on page 463. Once you have configured the LMP peer, you can obtain the traffic engineering link local ID and interface local ID by issuing the `show link-management te-link` command. Once you have these IDs, you can configure them as the remote IDs on the peer node.

To configure the remote ID for a traffic engineering link and for the traffic engineering link interface, include the `remote-id` statement:

```
te-link te-link-name {
  interface interface-name {
    remote-id id-number;
  }
  remote-id id-number;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring LMP Peers

You need to configure network peers for GMPLS. A peer is a network device that your router communicates with when setting up the control and data channels. The peer is often an optical cross-connect (OXC).

To configure an LMP peer name, include the `peer` statement at the [edit protocols link-management] hierarchy level:

```
peer peer-name {
  address ip-address;
  control-channel control-channel-interface;
  lmp-control-channel control-channel-interface {
    remote-address ip-address;
  }
  lmp-protocol {
    hello-interval milliseconds;
    hello-dead-interval milliseconds;
    retransmission-interval milliseconds;
    retry-limit number;
    passive;
  }
  te-link te-link-name;
}
```

The following sections describe how to configure the other statements needed for an LMP peer:

- Configuring the LMP Peer ID on page 463
- Configuring the Control Channel Interface on page 464
- Configuring the LMP Control Channel Interface for the Peer on page 464
- Configuring the Remote IP Address for the LMP Control Channel on page 465
- Configuring the Hello Message Attributes for the LMP Control Channel on page 465
- Configuring Message Attributes for the LMP Control Channel on page 466
- Configuring the Local Peer to Wait for the Remote Peer on page 467
- Configuring the Traffic Engineering Link for the LMP Peer on page 467
- Disabling the Traffic Engineering Link for the LMP Peer on page 467

Configuring the LMP Peer ID

To configure the LMP peer ID, include the `address` statement at the [edit protocols link-management peer *peer-name*] hierarchy level. The default value for the LMP peer ID is the loopback address.

```
[edit protocols link-management peer peer-name]
address ip-address;
```

Configuring the Control Channel Interface

You must configure one or more control channels between the LMP peers. The control channels must travel across either a point-to-point link or a tunnel.

To configure an interface for the control channel, include the `control-channel` statement at the `[edit protocols link-management peer peer-name]` hierarchy level:

```
[edit protocols link-management peer peer-name]  
  control-channel control-channel-interface;
```

You can configure a generic routing encapsulation (GRE) interface for the control channel. This type of interface does not require a Tunnel PIC.



NOTE: You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring the LMP Control Channel Interface for the Peer

In an environment that uses LMP to establish and maintain an LMP control channel between peers, you can configure a number of attributes associated with LMP. To configure the interface to be associated with the LMP control channel for the peer, include the `lmp-control-channel` statement:

```
lmp-control-channel control-channel-interface;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols link-management peer peer-name]`
- `[edit logical-routers logical-router-name protocols link-management peer peer-name]`

You can configure a generic routing encapsulation (GRE) interface for the LMP control channel. This type of interface does not require a Tunnel PIC.



NOTE: You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

When this LMP control channel interface comes up, the peers use LMP to negotiate channel parameters and configure the control channel.

The local peer repeatedly sends a Config message to the remote peer. The Config message contains the local control channel ID, the local peer's node ID, a message ID, and a CONFIG object that includes hello message attributes (the hello interval and the hello dead interval).

The channel is activated when the remote peer responds with a ConfigAck message. The remote peer does so only when its own configured hello interval and hello dead interval match the values in the received Config message or the default values. If these values do not match, the remote peer responds with a ConfigNack message. The local peer logs this event and resends the Config message until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the configuration process.

Configuring the Remote IP Address for the LMP Control Channel

You need to specify the address of the remote end of the LMP control channel.

To configure the remote IP address for the LMP control channel, include the `remote-address` statement:

```
remote-address address;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-control-channel *control-channel-interface*]
- [edit logical-routers *logical-router-name* protocols link-management peer *peer-name* lmp-control-channel *control-channel-interface*]

Configuring the Hello Message Attributes for the LMP Control Channel

Hello messages are exchanged between LMP peers to maintain the control channel after LMP has activated the control channel. The LMP control channel is considered to be up only when the hello negotiation is successful. Successful negotiation consists of the local peer sending a hello message to the remote peer and receiving a hello message in response.

The LMP peers continue to exchange hello messages after the LMP control channel is up in order to maintain the channel.

The hello interval specifies the interval between periodic hello messages. The hello dead interval specifies how long the local peer waits for a hello response before it declares the LMP control channel to be down. When the channel goes down, the local peer restarts the LMP control channel negotiation and configuration process.

You can specify a hello interval from 150 through 300,000 milliseconds. The default hello interval is 150 milliseconds.

You can specify a hello dead interval from 500 through 300,000 milliseconds. The default hello dead interval is 500 milliseconds.

To configure the attributes for hello messages exchanged between LMP peers, include the `hello-interval` and `hello-dead-interval` statements:

```
hello-interval milliseconds;  
hello-dead-interval milliseconds;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-routers *logical-router-name* protocols link-management peer *peer-name* lmp-protocol]

When an LMP control channel comes up after a successful exchange of hello messages between LMP peers, LMP uses link property correlation to verify the traffic engineering and data link information on both sides of a link. To do so, the local peer sends a LinkSummary message for each traffic engineering link governed by the LMP control channel. The LinkSummary message contains information that characterizes the traffic engineering link and each data link in the traffic engineering link.

The local peer continues sending a LinkSummary message for each link until the remote peer responds with a LinkSummaryAck message or until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the link property correlation process.

When the remote peer receives a LinkSummary message, it examines its own link information. If this information agrees with that in the LinkSummary message, the remote peer responds with a LinkSummaryAck message. If the information is different, the remote peer responds with a LinkSummaryNack message.

Configuring Message Attributes for the LMP Control Channel

You can configure message attributes that control the exchange of LMP Config and LinkSummary messages. The retransmission interval specifies the interval between resubmitted LMP messages. The retry limit specifies how many times LMP sends a message before restarting the process.

You can specify a retransmission interval from 500 through 300,000 milliseconds. The default retransmission interval is 500 milliseconds.

You can specify a retry limit from 3 through 1000 attempts. The default number of retry attempts is three.

To configure attributes governing the exchange of LMP messages between peers, include the `retransmission-interval` and `retry-limit` statements:

```
retransmission-interval milliseconds;  
retry-limit number;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-routers *logical-router-name* protocols link-management peer *peer-name* lmp-protocol]

Configuring the Local Peer to Wait for the Remote Peer

You can specify that the local peer does not initiate LMP negotiation. Instead, the local peer waits for the remote peer to configure the LMP control channel.

To configure the local peer to wait for the remote peer to configure the LMP control channel, include the `passive` statement:

```
passive;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-routers *logical-router-name* protocols link-management peer *peer-name* lmp-protocol]

Configuring the Traffic Engineering Link for the LMP Peer

To specify the name of a traffic engineering link to be associated with this peer, include the `te-link` statement at the [edit protocols link-management peer *peer-name*] hierarchy level:

```
[edit protocols link-management peer peer-name]  
te-link te-link-name;
```

For information on how to configure a traffic engineering link, see “Configuring LMP Traffic Engineering Links” on page 461.

Disabling the Traffic Engineering Link for the LMP Peer

To disable a specific traffic engineering link, include the `disable` statement:

```
disable;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management te-link *te-link-name*]
- [edit logical-routers *logical-router-name* protocols link-management te-link *te-link-name*]

Configuring Peer Interfaces in RSVP and OSPF

After you have configured the LMP peers, add the peer interfaces to RSVP and OSPF. The peer interface name must match the peer name configured in LMP. Once the peer interfaces are added to the protocols, the traffic engineering link local and remote addresses can be signaled and advertised to peers like any other interface enabled for RSVP and OSPF. These addresses act as virtual interfaces for GMPLS.



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

To configure peer interfaces in RSVP and OSPF, complete the procedures in the following sections:

- Configuring Peer Interfaces in RSVP on page 468
- Configuring Peer Interfaces in OSPF on page 468
- Configuring the Hello Interval for Peer Interfaces on page 468

Configuring Peer Interfaces in RSVP

To configure RSVP signaling for LMP peers, configure the LMP peer interface by including the `peer-interface` statement at the `[edit protocols rsvp]` hierarchy level:

```
[edit protocols rsvp]
peer-interface peer-interface-name {
  (aggregate | no-aggregate);
  authentication-key key;
  disable;
  hello-interval seconds;
  (reliable | no-reliable);
}
```

The statements configured at the `[edit protocols rsvp peer-interface peer-interface-name]` hierarchy level have the same functionality as the statements configured at the `[edit protocols rsvp interface interface-name]` hierarchy level.

Configuring Peer Interfaces in OSPF

To configure OSPF routing for LMP peers, configure the name of the LMP peer by including the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level:

```
[edit protocols ospf area area-number]
peer-interface peer-interface-name {
  dead-interval seconds;
  disable;
  hello-interval seconds;
  retransmit-interval seconds;
  transit-delay seconds;
}
```

For information on how to configure OSPF statements, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring the Hello Interval for Peer Interfaces

Hello packets are used to indicate to neighboring routers that the peer interface is still up and running. The hello interval must be the same for all routers on a shared logical IP network. You can specify a hello interval from 1 through 255 seconds. The default hello interval is normally 10 seconds. For nonbroadcast networks, the default hello interval is 120 seconds.

To specify how often the router sends hello packets out the peer interface, configure the `hello-interval` statement:

```
hello-interval seconds;
```

You can configure this statement at the following hierarchy levels:

- [edit logical-routers *logical-router-name* protocols ospf area *area-number* peer-interface *peer-interface-name*]
- [edit protocols ospf area *area-number* peer-interface *peer-interface-name*]

Configuring MPLS Paths for GMPLS

As part of the configuration for GMPLS, you need to establish a Multiprotocol Label Switching (MPLS) path for each unique device connected through GMPLS. Configure the traffic engineering link remote address as the address at the [edit protocols mpls path *path-name*] hierarchy level. Constrained Shortest Path First (CSPF) is supported so you can choose either the **strict** or **loose** option with the address.

See “Configuring LMP” on page 460 for information about how to obtain a traffic engineering link remote address.

To configure the MPLS path, include the **path** statement at the [edit protocols mpls] hierarchy level:

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

For information about how to configure MPLS paths, see “Creating a Named Path” on page 66.

Tracing LMP Traffic

To trace LMP protocol traffic, include the `traceoptions` statement at the `[edit protocols link-management]` hierarchy level:

```
[edit protocols link-management]
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size>
    <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Use the `file` statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`.

The following trace flags display the operations associated with the sending and receiving of various LMP messages:

- `all`—Trace all available operations
- `hello-packets`—Trace hello packets on any LMP control channel
- `init`—Output from the initialization messages
- `packets`—Trace all packets other than hello packets on any LMP control channel
- `parse`—Operation of the parser
- `process`—Operation of the general configuration
- `route-socket`—Operation of route socket events
- `routing`—Operation of the routing protocols
- `server`—Server processing operations
- `show`—Servicing operations for `show` commands
- `state`—Trace state transitions of the LMP control channels and traffic engineering links

Each flag can carry one or more of the following flag modifiers:

- `detail`—Provide detailed trace information
- `receive`—Packets being received
- `send`—Packets being transmitted

Configuring MPLS LSPs for GMPLS

To enable the proper GMPLS switching parameters, configure the label-switched path (LSP) attributes that are appropriate for your network connection. The default value for `switching-type` is `psc-1`, which is also appropriate for standard MPLS.

To configure the LSP attributes, include the `lsp-attributes` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name]  
lsp-attributes {  
    encoding-type type;  
    gpid gpid;  
    signal-bandwidth type;  
    switching-type type;  
}
```

If you include the `no-cspf` statement in the label-switched path configuration, you must also configure primary and secondary paths, or the configuration cannot be committed.

The following sections describe how to configure each of the LSP attributes for a GMPLS LSP:

- Configuring the Encoding Type on page 472
- Configuring the GPID on page 472
- Configuring the Signal Bandwidth Type on page 473
- Configuring GMPLS Bidirectional LSPs on page 473

Configuring the Encoding Type

You need to specify the encoding type of the payload carried by the LSP. It can be any of the following:

- ethernet—Ethernet
- packet—Packet
- pdh—plesiochronous digital hierarchy (PDH)
- sonet-sdh—SONET/SDH

The default value is `packet`.

To configure the encoding type, include the `encoding-type` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  encoding-type type;
```

Configuring the GPID

You need to specify the type of payload carried by the LSP. The payload is the type of packet underneath the MPLS label. The payload is specified by the generalized payload identifier (GPID).

You can specify the GPID with any of the following values:

- hdlc—High-Level Data Link Control (HDLC)
- ethernet—Ethernet
- ipv4—Internet Protocol version 4 (default)
- pos-scrambling-crc-16—For interoperability with other vendors' equipment
- pos-no-scrambling-crc-16—For interoperability with other vendors' equipment
- pos-scrambling-crc-32—For interoperability with other vendors' equipment
- pos-no-scrambling-crc-32—For interoperability with other vendors' equipment
- ppp—Point-to-Point Protocol (PPP)

To configure the GPID, include the `gpid` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  gpid gpid;
```

Configuring the Signal Bandwidth Type

The signal bandwidth type is the encoding used for path computation and admission control. To configure the signal bandwidth type, include the `signal-bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  signal-bandwidth type;
```

Configuring GMPLS Bidirectional LSPs

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, whereas GMPLS nonpacket 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`. Include the `switching-type` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  switching-type (lambda | fiber | ethernet);
```

Gracefully Tearing Down GMPLS LSPs

You can gracefully tear down nonpacket GMPLS LSPs. An LSP that is torn down abruptly, a common process in a packet-switched network, can cause stability problems in non-packet-switched networks. To maintain the stability of non-packet-switched networks, it might be necessary to tear down LSPs gracefully.

The following sections describe how to tear down GMPLS LSPs gracefully:

- Temporarily Deleting a GMPLS LSP on page 474
- Permanently Deleting a GMPLS LSP on page 474
- Configuring the Graceful Deletion Timeout Interval on page 475

Temporarily Deleting a GMPLS LSP

You can gracefully tear down a GMPLS LSP using the `clear rsvp session gracefully` command.

This command gracefully tears down an RSVP session for a nonpacket LSP in two passes. In the first pass, the `Admin_Status` object is signaled along the path to the endpoint of the LSP. During the second pass, the LSP is taken down. Using this command, the LSP is taken down temporarily. After the appropriate interval, the GMPLS LSP is resignaled and then reestablished.

The `clear rsvp session gracefully` command has the following properties:

- It only works on the ingress and egress routers of the RSVP session. If used on a transit router, it has the same behavior as the `clear rsvp session` command.
- It only works for nonpacket LSPs. If used with packet LSPs, it has the same behavior as the `clear rsvp session` command.

For more information, see the *JUNOS Routing Protocols and Policies Command Reference*.

Permanently Deleting a GMPLS LSP

When you disable an LSP in the configuration, the LSP is permanently deleted. By configuring the `disable` statement, you can disable a GMPLS LSP permanently. If the LSP being disabled is a nonpacket LSP, then the graceful LSP tear-down procedures that use the `Admin_Status` object are used. If the LSP being disabled is a packet LSP, then the regular signaling procedures for LSP deletion are used.

To disable a GMPLS LSP, include the `disable` statement at any of the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name]`—Disables the LSP.
- `[edit protocols link-management te-link te-link-name]`—Disables a traffic engineering link.
- `[edit protocols link-management te-link te-link-name interface interface-name]`—Disables an interface used by a traffic engineering link.

Configuring the Graceful Deletion Timeout Interval

The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

The ingress router initiates the graceful deletion procedure by sending the Admin_Status object in the Path message with the D bit set. The ingress router expects to receive an Resv message with the D bit set from the egress router. If the ingress router does not receive this message within the time specified by the graceful deletion timeout interval, it initiates a forced tear-down of the LSP by sending a PathTear message.

To configure the graceful deletion timeout interval, include the `graceful-deletion-timeout` statement at the `[edit protocols rsvp]` hierarchy level. You can configure a time from between 1 through 300 seconds. The default value is 30 seconds.

```
graceful-deletion-timeout seconds;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols rsvp]`
- `[edit logical-routers logical-router-name protocols rsvp]`

You can use the `show rsvp version` command to determine the current value configured for the graceful deletion timeout.

