

## Configuring GMPLS with CLI

This section provides information to configure UNI GMPLS using the command line interface.

Topics in this section include:

- [GMPLS Configuration Overview on page 462](#)
- [LMP and IPCC Configuration on page 463](#)
  - ☞ [Configuration of IP Communication Channels for LMP and RSVP on page 463](#)
  - ☞ [Configuring LMP on page 463](#)
  - ☞ [Configuring Traffic Engineering Links and Data Bearers on page 465](#)
- [Configuring MPLS Paths for GMPLS on page 468](#)
- [Configuring RSVP in GMPLS on page 470](#)
- [Configuring a GMPLS LSP on the UNI on page 472](#)
  - ☞ [gLSP Constraints on page 474](#)
- [Bandwidth on page 475](#)
- [Shared Risk Link Groups on page 476](#)
- [Optical Network Segment Recovery on page 478](#)
- [Configuration of End-to-End GMPLS Recovery on page 479](#)
- [GMPLS Tunnel Groups on page 482](#)
- [Configuring IP and MPLS in an Overlay Network to Use a GMPLS LSP on page 484](#)
- [Configuration Notes on page 485](#)

## GMPLS Configuration Overview

The Generalized Multi-Protocol Label Switching (GMPLS) User to Network Interface (UNI) permits dynamic provisioning of optical transport connections between IP routers and optical network elements in order to reduce the operational time and administrative overhead required to provision new connectivity.

---

## LMP and IPCC Configuration

### Configuration of IP Communication Channels for LMP and RSVP

Configuration starts with enabling the IP communication channel (IPCC) between the 7x50 UNI-C and the adjacent UNI-N. The IPCC is a data communication channel for LMP and RSVP. For each different 7x50 and UNI-N adjacency, a different IPCC must be configured.

In release 13.0, a numbered network IP interface is bound to the port connected to the DCN or directly to the 1830 PSS.

GMPLS protocols use a new loopback address type, called a **gmpls-loopback**, on the IPCC. The address of this loopback is termed the local GMPLS router ID. Packets that do not belong to a GMPLS protocol that are destined for this loopback address will be dropped. An interface is configured as a GMPLS Loopback using the **gmpls-loopback** keyword.

```
config
router
  interface local-gmpls-router-id-name
    gmpls-loopback
    address local-gmpls-loopback-address //Local LmpNodeId
```

The destination address of the LMP and RSVP control plane packets should be set to the LMP/GMPLS loopback of the 1830 PSS. The 1830 PSS does that via a dedicated subnet on a VLAN interface on the management port. Another VLAN extends a separate subnet for management purposes. On the 7x50 LMP and RSVP control plane packets should be sent to the next-hop for the GMPLS/LMP loopback address of the neighboring 1830 PSS. This is achieved via a static route in Release 13.0. The 1830 PSS and 7x50 GMPLS Router IDs must be in the same subnet. It may be possible to operate over a routed DCN network if the RSVP control plane messages will not set the IP router alert bit. Otherwise only direct IP connectivity, via a L2 network, will work.

If the IPCC goes down, then an existing TE Link or gLSP to a given peer UNI-N node is not torn down just because the IPCC is down. However, if the IPCC is down, then it is not possible to establish new gLSPs or TE Links, and a trap indicating a degraded state is raised.

## Configuring LMP

LMP is used to establish and maintain an IPCC between adjacent peers, as well as to correlate the local and remote identifiers for the TE Links that it controls. Some attributes must be configured locally on a per-peer basis, such as the LMP peer information, te-link information, and per-peer protocol related parameters.

The **config>router>lmp>lmp-peer peer-cp-node-id** command creates a context per LMP peer. The entry *peer-cp-node-id* is the control plane identifier of the adjacent UNI-N. It is an IPv4 or unsigned integer-formatted address that is used by the UNI-C for LMP and RSVP-TE

messages if a peer-loopback address is not subsequently configured. The local GMPLS Router ID is used as the source address.

In Release 13.0, a static route must have previously been configured to this peer router ID. Dynamic routing e.g. using OSPF over the IPCC in order to resolve routes to the peer GMPLS router ID, is not supported. In addition, the local loopback address to use as the local GMPLS Router ID should also be configured.

The LMP messages are sent over the interface corresponding to the IPCC that has been configured previously. The LMP session can be associated with one or more TE links that have been configured previously.

A control channel to an LMP Peer is configured using the **config>router>lmp>lmp-peer peer-cp-node-id>control-channel** context. Control channels are indexed using the *lmp-cc-id* parameter, which corresponds to the *lmpCcId* object in the LMP MIB.

The following CLI tree illustrates the key commands for configuring LMP.

```

config
  router
    [no] lmp
      [no] te-link te-link-id
        link-name te-link-name
        remote-id id
        [no] data-bearer data-bearer-id
          port port-id
          remote-id id
          [no] shutdown
        [no] shutdown

      gmpls-loopback-address local-gmpls-loopback-address
      [no] lmp-peer peer-cp-node-id
        peer-loopback-address peer-loopback-address
        retransmission-interval interval
        retry-limit limit
        [no] control-channel lmp-cc-id
          peer-interface-address ipcc-destination-addr
          hello interval interval dead-interval interval
          passive
          [no] shutdown
          te-link te-link-id
          te-link te-link-id
          [no] shutdown
        lmp-peer lmp-peer-address
        ...
        [no] shutdown
      [no] shutdown
  
```

If **peer-loopback-address** is entered, then this is used as the routable peer address, otherwise the *peer-cp-node-id* is assumed to correspond to a routable peer loopback.

The **peer-interface-address** is mandatory and is the destination address of the IPCC on the peer UNI-N used to reach the GMPLS Router ID of the peer. It corresponds to the

ImpCcRemoteIpAddr in RFC 4631. If the **peer-interface-address** is used as the destination IP address in the IP packet on the IPCC, then the 7x50 local interface address is used as the source IP address.

A **te-link** is configured under **config>router>imp>te-link**. The **te-link** parameter under **config>router>imp>imp-peer** then assigns the control of the TE-links to the LMP protocol to a given peer. Each TE-Link can only be assigned to a single LMP peer.

The LMP protocol-specific attributes such as timers and retransmission retries are configured for each LMP peer under **configure>router>imp>imp-peer**.

The **hello interval** ranges from 1000 to 65535 milliseconds. The default hello interval is 1000 milliseconds.

The **hello dead-interval** ranges from 3000 to 65535 milliseconds. The default hello dead interval is 4000 milliseconds.

The **retransmission-interval** ranges from 10 to 4294967295 milliseconds in 10 millisecond intervals, with a default of 500 milliseconds.

## Configuring Traffic Engineering Links and Data Bearers

Traffic engineering (TE) links are configured under the **config>router>imp** with a specific command, **te-link**, to create a specific context to hold TE specific configuration information pertinent to the local and remote identifiers, and physical resources assigned to the te-link. Only one data bearer per TE link is supported.

The te-link association is the creation of an association between a TE-link and data-bearing physical ports. Under the TE-link context, different data bearers can be configured via the data-bearer command. The data bearer is assigned a complete physical port, using **port<x/y/z>** (slot-number/MDA-number/port-number) as input.

Note that a data bearer cannot be associated with a port in a LAG.

A TE-link has a unique *link-id*, which identifies it in RSVP-TE signaling.

The remote-id is the unnumbered link identifier at far-end of the TE link as advertised by the LMP peer i.e. the UNI-N.

The TE-link has associated physical resources which are assigned to the TE-link by configuring the data-bearer under the **config>router>te-link** context.

The operator must also configure the remote data-bearer link identifier under the data bearer subcontext.

Note that LMP does not correlate the local and remote Layer 2 interface identifiers (such as MAC addresses) for the data bearer. It only correlates the local and remote TE Link and Data Bearer link identifiers. The association between the Layer 2 interface address and the data bearer must be correctly configured at the UNI-C and UNI-N. The **show>router>lmp>te-link** command displays the local link ID, remote link ID, and associated port ID to assist with this.

The CLI tree for creating TE Links under LMP is as follows. Note that there are also some RSVP-specific TE Link parameters that are configured under a separate **gmpls** context (see below):

```
config
  router
    [no] lmp
      [no] te-link te-link-id
        link-name te-link-name
        remote-id id
        [no] data-bearer data-bearer-id
          port port-id
          remote-id id
          [no] shutdown
        [no] shutdown
      [no] shutdown
```

The *te-link-id* can take the form of an unsigned integer or 64 character (max) name: [1..2147483690] | *te-link-name*: 64 char max

Upon creation, only the unsigned integer needs to be specified. Once the link is created the user can configure the link name (ie. 'link-name *te-link-name*'). From here, the user can refer to this te-link by either the unsigned integer or the ASCII name.

Note that LMP will normally assume a data bearer is operationally up, even if no MAC layer or a valid PCS IDLE stream is received. This is because a neighboring UNI-N may not generate a PCS IDLE stream and instead transparently transports the MAC layer from the far end, which won't be up unless a gLSP is configured. In order to prevent LMP from using a port for which there is a local fault on the data bearer, indicated by loss of light, a user must configure **report-alarm** on the Ethernet port, as follows:

```
config>port>ethernet>report-alarm signal-fail
```

Only ports with **report-alarm signal-fail** configured can be included in LMP, and that **report-alarm signal-fail** cannot be subsequently removed from a port in LMP.

RSVP requires that all traffic engineering attributes for TE Links are configured under the **config>router>gmpls>te-link** context.

```
config
  router
    [no] gmpls
      te-link te-link-id
        [no] shutdown
```

where *te-link-id*: [1..2147483690] | *te-link-name*: 32 char max

If a path (also refer to the description of a GMPLS path configuration, below) without an explicit *te-link* for the first hop is configured, the system will automatically select a TE Link to use for a gLSP path based on the lowest available TE Link ID with a matching bandwidth (if a bandwidth is configured for the gLSP). During a data-bearer link allocation request, an RSVP-requested gLSP BW could be either a non-zero value as per RFC 3471 signal-type (see below), or it could be zero. There are the following cases

#### **Case 1: Requested BW is non-zero as per RFC 3471 Signal-type configuration**

- When a TE (or TE/DB) link is configured in the related hop LMP checks whether the related port BW is the same (exact match) as the requested BW, and allocates the port (provided any other checks are successful).
- When the related Hop is empty, LMP finds a db-link port to the peer with a matching the requested BW, and allocates it.

#### **Case 2: Requested BW is Zero**

- When TE (or TE/DB) link is configured in the related hop, LMP allocates the port (provided the other checks are OK), and provides the port BW to RSVP to use in signaling.
- When the related Hop is empty, LMP finds the first available db-link to the peer (based on lower db-link Id), and allocates it and provides the port BW to RSVP to use in signaling.

## Configuring MPLS Paths for GMPLS

To establish an end-to-end connection between two 7x50s through a GMPLS network, a path is required, which is configured via the **configure>router>gmpls>path** *path-name* context.

The path context consists of a set of numbered entries, each entry representing a resource that the gLSP must follow. Note that the te-link ID is the ID allocated at the node referred to in the hop.

When interoperating with the 1830 PSS, at least the first and penultimate hops of the gLSP should be included.

The following CLI tree is used to configure a gLSP path:

```
config>router>gmpls
  path path-name
  no path path-name
  hop hop-index node-id node-id [te-link te-link-id]
    [strict | loose]
  no hop hop-index
  no shutdown
  shutdown
```

where:

*node-id*: IPv4 address a.b.c.d | 1830-data-plane-node-id 32-bit unsigned integer

In general, the 7x50 is able to populate the ERO with every hop along the gLSP path from ingress UNI-N to egress UNI-C. However, normally only a loose path across the optical network (from ingress UNI-N to egress UNI-N) is required because the optical network is responsible for path selection between ingress and egress UNI-N. Therefore the user will normally just configure hop 1 and hop 4 in the above example. For interoperability with the 1830 PSS, the user must configure a TE Link ID to use on the final hop in the ERO towards the destination UNI-C.

The following example shows how the Path should be configured for interoperability with the 1830 PSS.

Consider the following topology:

```

      A      B      C      D      E      F
[unic1]-----[unin1]-----[unin2]-----[unic2]
```

where A-F are the TE Link IDs assigned at each end of a link.

Path configuration on unic1:

```
Hop 1 unic1 A strict
```



Hop 2 unin2 E loose

## Configuring RSVP in GMPLS

RSVP-TE must be enabled on the SR OS towards the adjacent UNI-N in order to configure a GMPLS label-switched path (gLSP).

RSVP parameters specific to GMPLS are configured under the **config>router>gmpls** context.

This creates a new instance of RSVP for use in GMPLS signaling.

Global parameters for GMPLS are configured as follows:

```
config
  router
    gmpls
  no gmpls
    peer peer-cp-node-id
    gr-helper-time max-recovery recovery-interval max-restart restart-interval
  no gr-helper-time
    keep-multiplier number
  no keep-multiplier
  no rapid-retransmit-time
  rapid-retransmit-time hundred-milliseconds
  no rapid-retry-limit
  rapid-retry-limit limit
  no refresh-time
  refresh-time seconds
  no refresh-time
  lsp-init-retry-timeout seconds
  no lsp-init-retry-timeout
  no shutdown
  shutdown
```

The default max-restart interval for GMPLS is 180 seconds.

The LMP Peer is configured under **config>router>gmpls>peer** *peer-cp-node-id*, where the *peer-cp-node-id* is control plane identifier of the adjacent optical cross connect (UNI-N node). RSVP uses the destination address returned by LMP for this peer control plane node ID as the destination address, and the loopback address referenced under **config>router>lmp>gmpls-loopback-address** *local-gmpls-loopback-address* as the local router ID to use for the session.

RSVP will come up if at least one IPCC is up.

RSVP hellos and support for graceful restart helper functionality are supported. RSVP Graceful Restart Helper procedures implemented by the 7x50 also apply when the IPCC goes down and comes back up, or when the neighboring peer control plane restarts.

The following CLI tree is used for configuring RSVP parameters for each LMP peer:

```
config
  router
    gmpls
```

```
peer peer-cp-node-id
no peer peer-cp-node-id
  lsp-hold-timer hold-timer
  no lsp-hold-timer
  hello-interval milliseconds
  no shutdown
  shutdown
```

The per-peer **lsp-hold-timer** *hold-timer* parameter is used to configure a node-wide hold-down time. This timer is started when a RESV for a new gLSP is first received, or a failed gLSP path is restored (or the 7x50 is notified of a restoration following segment recovery) in order to give the optical network time to program its data path. The value range is 5 to 300 seconds, with a default of 60 seconds. A member of a GMPLS tunnel group is not considered up until the hold-timer has expired. Note that different optical network technologies have different data path programming/setup times.

Note that the **no hello-interval** command sets the hello-interval to the default value of 3000 milliseconds. Configuring **hello-interval 0** will disable hellos in GMPLS.

## Configuring a GMPLS LSP on the UNI

A GMPLS LSP is configured under `config>router>gmpls>lsp name gmpls-uni`. The optional `gmpls-uni` keyword indicates that the LSP is an RSVP signaled GMPLS LSP, which is profiled for the GMPLS UNI i.e. it uses the set of functions and CLI commands applicable to an overlay gLSP, rather than a peer model gLSP. Only overlay model gLSPs are supported in Release 13.0; this is the default type of GMPLS LSP. The 7x50 can only act as an LER terminating a gLSP, and cannot switch a GMPLS i.e. it cannot act as a GMPLS LSR

GMPLS LSPs use the working path and protect path terminology from RFC 4872. Each gLSP configuration is composed of a working path and an optional protect path if end-to-end recovery is used.

Note that on-the-fly changes to an LSP or LSP path configuration are not allowed. This is because MBB is not supported for gLSPs. The LSP or LSP Path must be shut down to make configuration changes.

A GMPLS LSP (gLSP) is configured using the following CLI tree:

```

config
  router
    gmpls
      lsp lsp-name [gmpls-uni]
      no lsp lsp-name
        to remote-uni-c-gmpls-router-id
        switching-type {dcsc}
        no switching-type
        encoding-type {line}
        no encoding-type
        generalized-pid {ethernet}
        no generalized-pid
        e2e-protection-type {unprotected|1toN | sbr}
        no e2e-protection-type
        protect-path path-name
        no protect-path path-name
          peer peer-gmpls-router-id
          no peer
          bandwidth signal-type rfc3471-name
          no bandwidth exclude-srlg group-name [group-name...(upto 5 max)]
          no exclude-srlg
          segment-protection-type {unprotected|sbr|gr|sncp|prc}
          no segment-protection-type
          no shutdown
          shutdown
          revert-timer timer-value //1 to 1800 seconds, default 0
          no revert-timer
          retry-limit limit
          no retry-limit
          no shutdown
          shutdown
          working-path path-name
          no working-path path-name
            bandwidth signal-type rfc3471-name
  
```

```

no bandwidth
exclude-srlg group-name [group-name... (upto 5 max)]
no exclude-srlg
peer peer-gmpls-router-id
no peer bandwidth
segment-protection-type {unprotected | sbr | gr | sncp | prc}
no segment-protection-type
no shutdown
shutdown
no shutdown
shutdown

```

The loopback address of the remote 7x50 (UNI-C) must be configured after the **to** keyword and takes an IPv4 address as input.

The **switching-type** indicates the type of switching required for the gLSP. This can take a number of values, as defined in RFC 3471, and extended in RFC 6004 and RFC 7074 for Ethernet VPL (EVPL) services. The default CLI value is **DCSC**. This is the only supported value in Release 13.0.

The **encoding-type** configuration specifies the encoding type of the payload carried by the gLSP. **line**, indicating 8B/10B encoding, is the only supported type in Release 13.0.

The **generalized-pid** parameter specifies the type of payload carried by the gLSP. Standard ethertype values are used for packet and Ethernet LSPs (see RFC 3471). Only Ethernet (value 33) is supported in Release 13.0.

Note that gLSPs are inherently bidirectional. That is, both directions of the gLSP are bound together. The destination UNI-C 7x50 will automatically bind an incoming gLSP PATH message to the corresponding egress direction based on the session name in the session object.

Any gLSP that needs to be bound to a specific TE Link (as referred to in the pPATH), will only be allowed if the corresponding TE Link exists under `config>router>gmpls`. Constraints such as HOP definition, SRLG, BW, etc., will be checked before signaling the gLSP.

Since RSVP signaling operates out of band, refresh reduction is not supported. RSVP authentication is not supported on the 1830 UNI-N, but MD5 authentication is implemented.

A configurable **retry-timer** is supported.

A configurable **retry-limit** for each gLSP is supported, with a range of 0 to 10000, and a default of 0.

The **working-path** and **protect-path** command allows paths to be configured for the gLSP. At least a **working-path** must be configured, although the path-name that it references may contain an empty path. The optional **working-path>peer** and **protect-path>peer** commands allow the user to specify a first hop UNI-N node to use for the gLSP path. The protect path is only configurable for 1:N recovery option.

Reversion from the protect path to the working path is supported.

RSVP uses the Fixed Filter (FF) style of RESV. The signaled MTU is hard-coded to 9212 bytes, as appropriate for Ethernet gLSPs.

The default **setup** and **hold** priorities are 5 and 1, respectively, and cannot be configured in Release 13.0. gLSP preemption is not supported.

**Record** and **record-label** are enabled by default and no user configurable command is therefore provided.

---

## gLSP Constraints

Each gLSP can be configured with the following constraints:

- Bandwidth
- SRLG
- Protection

## Bandwidth

The bandwidth associated with a gLSP is configured with the bandwidth command, and can take the RFC 3471 signal type name as input in Release 13.0.

The signaled bandwidth is then used for path computation and admission in the GMPLS domain.

By default the, actual interface bandwidth is used. If the user configures a bandwidth greater than the local data bearer bandwidth, then the gLSP establishment will be blocked. If the user configures a bandwidth less than or equal to the local data bearer bandwidth, then that bandwidth is signaled to the UNI-N.

The bandwidth required for the LSP is configured under the path context as follows. Note that the system will do an exact match check of the gLSP bandwidth against the data bearer bandwidth:

```
config
  router
    gmpls
      lsp gmpls-tunnel-name [gmpls-uni]
        to remote-uni-c-gmpls-router-id
        working-path path-name
        bandwidth signal-type rfc3471-name
```

The possible signal-type values are:

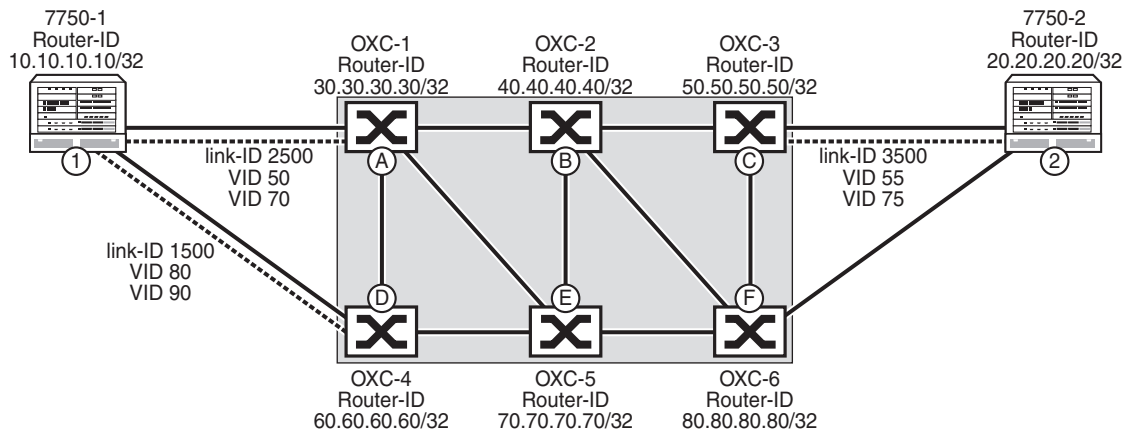
ds0 | ds1 | e1 | ds2 | e2 | ethernet | e3 | ds3 | sts-1 | fast-ethernet | e4 | fc-0-133m | oc-3/stm-1 | fc-0-266m | fc-0-531m | oc-12/stm-4 | gige | fc-0-1062m | oc-48/stm-16 | oc-192/stm-64 | 10gige-ieee | oc-768/stm-256 | 100gige-ieee

The code points to use for 10gige-ieee and 100gige-ieee are not yet registered with IANA. The following values are therefore used:

- 10G IEEE: 0x4E9502F9
- 100G IEEE: 0x503A43B7

## Shared Risk Link Groups

Shared Risk Link Groups (SRLG) are used in the context of a gLSP to ensure that diverse paths can be taken for different gLSPs through the optical network. For example, consider the network shown in the following figure:



24854

**Figure 48: SRLG Example**

In this dual-homing scenario, the primary gLSP takes TE-Link 1-A, and C-2, while the secondary gLSP path takes TE-Links 1-D and F-2. In order to ensure that a failure in the underlying optical network does not affect both the primary and secondary paths for the gLSP, the SRLG list used by the optical network for the primary path is shared with the UNI-C (1) by the UNI-N (A) at the time the gLSP is established along the primary path. When the secondary path is signaled, the UNI-C (1) will signal the SRLG list to avoid to the UNI-N (D). Note that a similar procedure is beneficial even if a UNI-C is not dual homed to the optical network, but diverse primary and secondary paths are required through the optical network.

The 7x50 supports two methods for indicating a set of SRLGs to exclude:

- Explicit configuration of an SRLG list for a gLSP path. These are signaled in the XRO of the RSVP PATH message towards the optical network
- Automatic SRLG collection for a gLSP, using the procedures specified in draft-ietf-ccamp-rsvp-te-srlg-collect-04.txt, and operate as follows:
  - ☞ Retrieving SRLG information from a UNI-N for an existing gLSP Path — When a dual-homed UNI-C device intends to establish a gLSP path to the same destination UNI-N device via another UNI-N node, it can request the SRLG information for an already established gLSP path by setting the SRLG information flag in the LSP attributes sub-object of the RSVP PATH message



using a new SRLG flag. This path would be the primary path for a gLSP established by the 7x50 UNI-C. As long as the SRLG information flag is set in the PATH message, the UNI-N node inserts the SRLG sub-object as defined in draft-ietf-ccamp-rsvp-te-srlg-collect-04.txt into the RSVP RESV message that contains the current SRLG information for the gLSP path. Note that the provider network's policy may have been configured so as not to share SRLG information with the client network. In this case the SRLG sub-object is not inserted in the RESV message even if the SRLG information flag was set in the received PATH message. Note that the SRLG information is assumed to be always up-to-date by the UNI-C.

- ☞ Establishment of a new gLSP path with SRLG diversity constraints — When a dual-homed UNI-C device sends an LSP setup requests to a UNI-N for a new gLSP path that is required to be SRLG diverse with respect to an existing gLSP path that is entering the optical network via another UNI-N, the UNI-C sets a new SRLG diversity flag in the LSP attributes sub-object of the PATH message that initiates the setup of this new gLSP path. This path would be the protect path of a gLSP established by the 7x50. When the UNI-N receives this request it calculates a path to the given destination and uses the received SRLG information as path computation constraints.

In Release 13.0, the 7x50 collects SRLG by default. SRLG collection occurs on all paths of the gLSP. The collected SRLG list is visible to the user via a **show** command. The recorded SRLGs are then used to populate the XRO. Only best effort (ie. loose) SRLG diversity is supported.

Automated SRLG diversity is supported for the working and protect paths of the following end to end protection types in Release 13.0R1:

- 1:N
- LSPs that form a part of a load sharing tunnel group

Already-established gLSPs within a load-sharing tunnel group or for which 1:N recovery is configured can be made mutually diverse by applying a **shutdown / no shutdown** operation. GMPLS LSPs with other types of protection can be made mutually SRLG-diverse by performing a shutdown of the gLSP, reconfiguring the SLG list to exclude using the **exclude-srlg** command, and then applying a **no shutdown** of the gLSP.

## Optical Network Segment Recovery

The 7x50 may request a particular GMPLS recovery type for a gLSP path segment that spans the optical network. This refers to the protection afforded to the gLSP path between the UNI-N nodes. The 7x50 supports the following segment protection types (code points are also shown):

- Unprotected: 0x00
- Source-Based Reroute (SBR) (Known as Full Rerouting in the IETF): 0x01
- Guaranteed Restoration (GR) (Also known as shared mesh restoration): 0x02
- Sub-network Connection Protection (SNCP) (1+1 bidirectional protection): 0x10
- Path Restoration Combined (PRC): 0x11

These resiliency options are configured under the **segment-protection-type** command for a given path.

```
config
router
  gmpls
    lsp gmpls-tunnel-name [gmpls-uni]
      to remote-uni-c-gmpls-router-id
      working-path path-name
        [no] segment-protection-type {unprotected | sbr | gr | sncp | prc}
        ...
        [no] shutdown
```

The default **segment-protection-type** setting is **unprotected**.

If the requested protection type cannot be satisfied by the optical network, the 7x50 will generate a CLI warning and an SNMP trap.

The [Table 11](#) illustrates the recommended combinations of segment protection type and end-to-end protection type.

**Table 11: Combinations of End-to-End and Segment Protection**

E2E/Segment	Unprotected	SBR	GR	SNCP	PRC
Unprotected	Yes	Yes	Yes	Yes	Yes
1:1/1:N	Yes	Yes	Yes	Yes	No
Full Rerouting	Yes	No	No	Yes	No

## Configuration of End-to-End GMPLS Recovery

End-to-end GMPLS recovery is configured at the LSP level using the **e2e-protection-type** command, as follows:

```
config
  router
    gmpls
      lsp gmpls-tunnel-name [gmpls-uni]
        to remote-uni-c-gmpls-router-id
        e2e-protection-type [unprotected|1toN|sbr]
        revert-timer timer-value
```

The protection type names are common to those used in the optical network. The protection types are as follows:

- **unprotected** — 0x00
- **1toN** — 1:N protection. Extra traffic is not supported. Note that 1:1 protection is a special case of 1:N. 0x04
- **sbr** — Full LSP rerouting; 0x01

The default end-to-end protection type is **unprotected**.

It is possible to configure segment protection on a path independently of the type of end-to-end protection that is configured.

1toN protection requires the configuration of multiple working paths and a protect path for a GMPLS LSP. The working paths are then associated with different GMPLS Tunnel Groups. Configuration is as follows:

```
config
  router
    gmpls
      lsp lsp-name gmpls-uni
        to remote-uni-c-gmpls-router-id
        e2e-protection-type 1toN // Only these types are allowed for gmpls-uni
        switching-type ethernet
        encoding-type ethernet
        generalized-pid ethernet
        revert-timer timer-value
        retry-limit limit
        working-path path-name1 [lmp-peer <peer-gmpls-router-id>] ...
          [no] shutdown
        working-path path-name2 [lmp-peer peer-gmpls-router-id] ...
          [no] shutdown
        working-path path-name3 [lmp-peer peer-gmpls-router-id] ...
          [no] shutdown
        protect-path path-name4 [lmp-peer peer-gmpls-router-id] ...
          [no] shutdown
```

The LSP is then bound to one or more GMPLS tunnel groups. Load sharing or 1:N protection may be used across the working paths. The load sharing case is described below.

For the non-load sharing 1:N case, a single LSP is assigned to each tunnel group, as follows:

For the head end node:

```
config > gmpls-tunnel-group 2 create
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 3
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 4
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
```

For the tail end node:

```
config > gmpls-tunnel-group 2
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 3
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 4
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
```

```
no shutdown  
no shutdown
```

Note that a shutdown of a working path does not trigger a switchover to the protect path. The user should either use the `tools>perform>router>gmpls force` or manual commands, or shutdown the TE-Link, data bearer, or port associated with the gLSP path.

## GMPLS Tunnel Groups

A GMPLS tunnel group is a bundle of gLSPs providing an abstraction of the data bearers that are intended to be associated to one IP interface. This object allows, for example, end-to-end load balancing across the set of data bearers corresponding to a set of gLSPs. A gLSP is bound to a GMPLS tunnel group by a gLSP tunnel (session) name at both the head end and the tail end UNI-C nodes of the gLSP. A sender address (the far-end) may optionally be configured for the tail end of a gLSP in case different head end nodes use overlapping gLSP tunnel names.

```

config
  gmpls-tun-grp gmpls-tun-grp-id
    type {head-end | tail-end}
    far-end remote-uni-c-router-id
    mode {load-sharing | active-standby}
    no mode
    [no] member-threshold threshold [action down]
    member mem-id [create]
      glsp session-name name
      no glsp session-name name
      [no] shutdown
    ...
    [no] shutdown

```

*gmpls-tun-grp-id* is an unsigned integer from 1 to 1024, shared with the Ethernet tunnel ID range.

The GMPLS Tunnel Group must be configured as either at both the **head-end** or **tail-end** of a set of member gLSPs (identified using the **head-end** or **tail-end** keywords). These keywords are mutually exclusive.

Nodes at the head-end initiate signaling of gLSPs. The **far-end** is the far end of the GMPLS tunnel group. If this node is a head end, then the far end address is taken as the to address for the member gLSPs. Each gLSP that is bound to the tunnel group must have a to address matching the far end address. A binding is held down if a gLSP to and the tunnel group to do not match.

Nodes at the tail end wait for the first path message for a gLSP. The **far-end-address** address must be configured at the tail end. It is the GMPLS Router ID of the head-end UNI-C (the *remote-uni-c-node-id*), and must be configured at the tail end UNI-C of a gLSP. The combination of *session-name* and *remote-uni-c-node-id* provides a unique key to bind an incoming gLSP setup request to a tunnel group. A binding to the tunnel group is held down at the tail end until a gLSP PATH message with a matching *session-name* and source address that matches the tunnel group's far-end address is received.

At the tail end, the **session-name** is composed of the LSP name and Path name as configured at the head end

If **load-sharing** is configured, then all of the gLSPs must terminate on the same far-end node. All of the ports used by gLSPs in a load-sharing must be equivalent in that they must have the same named QoS policy, bandwidth, and so on. Once more than one gLSP is associated with a tunnel group, the QoS policy/scheduler policy cannot be changed in any of the ports. All gLSPs must be unprotected end-to-end in load-sharing mode. Segment protection is allowed for gLSPs associated in load sharing mode to a GMPLS tunnel group.

In **active-standby** mode, only one member gLSP can be associated with the tunnel group.

All members of a tunnel group must be of the same bandwidth.

The **member-threshold** is the number of member gLSPs that must be operationally up before the gmpls tunnel group is considered operationally up.

A member of a GMPLS tunnel group may be treated as down for one of the following reasons. These reason codes are recorded in the `tmnxGmplsTunGrpMemberTable` in the MIB:

- `adminDn` — The member or the related tunnel-grp is administratively down.
- `wpLspDn` — The associated working lsp-path is down.
- `wpPortDn` — The data-bearer port associated with the working lsp-path is down.
- `wpPortNoRsrc` — The data-bearer port associated with the working lsp-path has no resource to support the services over the gmpls-tunnel-grp logical port.
- `ppLspDn` — The associated protect lsp-path is down.
- `ppPortDn` — The data-bearer port associated with the protect lsp-path is down.
- `ppPortNoRsrc` — The data-bearer port associated with the protect lsp-path has no resource to support the services over the gmpls-tunnel-grp logical port.

Note that in the case of `wpPortNoRsrc` and `ppPortNoRsrc`, the term 'resources' relates to QoS or ACL related resources. For example, this can happen when a subsequent physical or data bearing port is added to a GMPLS tunnel group, which already has services running over it. If the new-complex doesn't have the resources to support those services over that GMPLS tunnel group, the related member `operState` would be down with reasonCode `PortNoRsrc`. If a gLSP is already established on a data bearer when a resource failure is experienced, the RSVP PATH message A-Bit is updated so that both ends ensure that the LSP Path is held down.

The user should free resources from the complex, and shutdown/no shutdown the GMPLS tunnel group member. This repeats the resource check, which will bring the member `operUp` if it passes.

A gLSP associated with a tunnel group member will be down if the member is operationally down, or a fault is detected on the associated data bearer.

If a member is in the admin down state, a gLSP will not be set-up. If a gLSP is already up, the RSVP Path message A-Bit is updated so that both ends of the gLSP path are kept down.

## Configuring IP and MPLS in an Overlay Network to Use a GMPLS LSP

IP and MPLS is able to use GMPLS LSPs as transport by bringing a numbered or unnumbered IP interface to an endpoint of one or more gLSPs. This IP interface appears as any other IP interface bound to a network port. The IP interface is bound to the GMPLS tunnel group by a GMPLS tunnel group number configured in the **port** command.

The GMPLS tunnel group number must correspond to a locally configured GMPLS tunnel group.

The following CLI tree illustrates where the GMPLS tunnel group is referenced. This must be done at nodes at 7x50 nodes at the tunnel groups at both ends of the transport service.

```
config
  router
    interface if-name
      address a.b.c.d|ipv6-address
      port gmpls-tunnel-group gmpls-tunnel-group-id
```



## Configuration Notes

This section describes GMPLS caveats.

- Interfaces must already be configured in the `config>router>interface` context before they can be specified in GMPLS.
- A router interface must be specified in the `config>router>mpls` context in order to apply it or modify parameters in the `config>router>rsvp` context.
- A system interface must be configured and specified in the `config>router>mpls` context.
- Paths must be created before they can be applied to an LSP.

