MPLS Transport Profile

In This Chapter

This section provides information about Multiprotocol Label Switching Transport Profile (MPLS-TP).

Topics in this section include:

Applicability

This example is applicable to the 7950, 7750 and 7450 series and was tested on release 12.0.R2. Multiprotocol Label Switching Transport Profile (MPLS-TP) requires a minimum of FP2 or higher based hardware. A CPM3 or higher is required for the highest BFD scale using 10ms control packet timers.

This example assumes that the reader is familiar with the configuration of IP/MPLS and VLL services on the 7x50.

MPLS-TP was first introduced in SR OS release 11.0.R4 and further enhancements were added in subsequent releases.

Summary

MPLS-TP is intended to allow MPLS to be operated in a similar manner to existing transport technologies, with static configuration of transport paths (particularly with no requirement for a dynamic control plane), in-band proactive and on-demand operations and maintenance (OAM), and protection mechanisms that do not rely on a control plane (for example, Resource Reservation Protocol with Traffic Engineering (RSVP-TE)) to operate. The 7x50 can operate both as a Label Edge Router (LER) and Label Switching Router (LSR) for MPLS-TP LSPs, and as a Terminating Provider Edge (T-PE) and Switching Provider Edge (S-PE) for Pseudowires (PWs) with MPLS-TP OAM. The 7x50 can therefore act as a node within an MPLS-TP network, or as a gateway between MPLS-TP and IP/MPLS domains.

Overview

MPLS can provide a network layer with packet transport services. In some operational environments it is desirable that the operation and maintenance of such an MPLS based packet transport network follows the operational models typically used in traditional optical transport networks (for example with SONET, SDH) while providing additional OAM, survivability and other maintenance functions targeted at that environment.

MPLS-TP defines a profile of MPLS targeted at transport applications. This profile defines the specific MPLS characteristics and extensions required to meet transport requirements, while retaining compliance with the standard IETF MPLS architecture and label-switching paradigm. The basic architecture and requirements for MPLS-TP are described by the IETF in RFC 5654, RFC 5921 and RFC 5960, in order to meet two objectives:

- To enable MPLS to be deployed in a transport network and operated in a similar manner to existing transport technologies.
- To enable MPLS to support packet transport services with a similar degree of predictability to that found in existing transport networks.

In order to meet these objectives, MPLS-TP has a number of high-level characteristics:

- MPLS-TP, including resilience and protection, operates in the absence of an IP control plane and IP. MPLS-TP does not modify the MPLS forwarding architecture, which is based on existing pseudowire and LSP constructs. Point-to-point LSPs may be unidirectional or bi-directional. Bi-directional LSPs must be congruent (i.e. co-routed and follow the same path in each direction) and are the only supported type on the 7x50. MPLS-TP is only supported on static LSPs and pseudowires (PWs). Also, there is no LSP merging.
- LSP and pseudowire monitoring is achieved using in-band OAM and does not rely on control plane or IP routing functions to determine the health of a path, for example, LDP hello failures do not trigger protection.

The system supports MPLS-TP on LSPs and PWs with static labels. MPLS-TP is not supported on dynamically signaled LSPs and PWs, although switching a static MPLS-TP PW to a targeted LDP (T-LDP) signaled PW is supported. MPLS-TP is supported for Epipe, Apipe and Cpipe VLLs, and Epipe spoke SDP termination on IES, VPRN and VPLS. Static PWs may use SDPs on top of either static MPLS-TP LSPs or RSVP-TE LSPs.

The following MPLS-TP OAM and protection mechanisms defined by the IETF are supported:

- MPLS-TP Generic Associated Channel for LSPs and PWs (RFC 5586)
- MPLS-TP Identifiers (RFC 6370)
- Proactive CC, CV, and RDI using BFD for LSPs (RFC 6428)
- On-Demand CV for LSPs and PWs using LSP Ping and LSP Trace (RFC 6426)

- 1-for-1 Linear protection for LSPs (RFC 6378)
- Static PW Status Signaling (RFC 6478)

The system can play the role of an LER and an LSR for static MPLS-TP LSPs, and a PE/T-PE and an S-PE for static MPLS-TP PWs. It can also act as an S-PE for MPLS-TP segments between an MPLS network that strictly follows the transport profile and an MPLS network that supports both MPLS-TP and dynamic IP/MPLS.

Configuration

This section details the configuration steps for a set of simple MPLS-TP examples.

The following reference network is used (Figure 162). It consists of four nodes and two Epipe VLL services. One service is transported across a network domain consisting of only static MPLS-TP LSPs (Epipe 10) from PE-1 to PE-2. The other Epipe (Epipe 20) is used to transport traffic from PE-1 in the MPLS-TP domain to a VPLS service on PE-4 in an IP/MPLS domain. A static MPLS-TP LSP exists between PE-1 and PE-2, while a dynamic RSVP-TE LSP exists between PE-2 and PE-4.

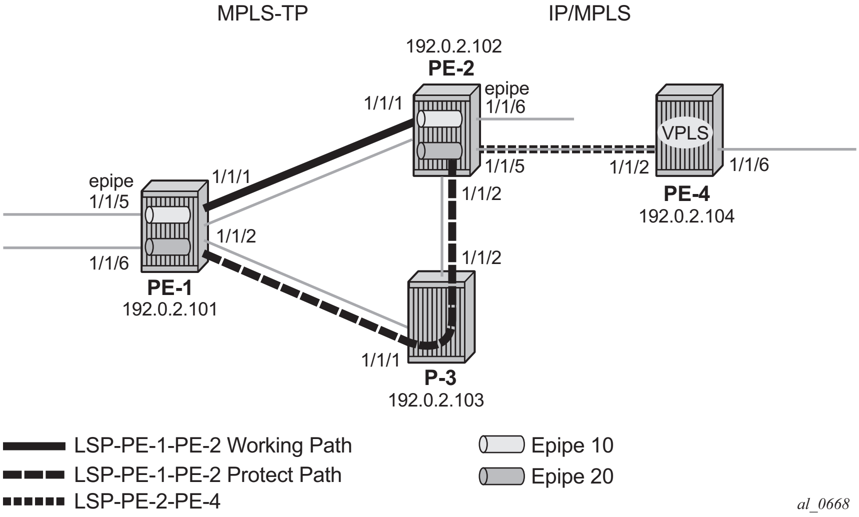

Figure 162: MPLS-TP Example Network Showing LSPs

Figure 163 shows further details of the logical architecture of the services in the example network. The Epipe spoke-sdps use the static MPLS-TP transport LSP between PE-1 and PE-2, and the dynamically signaled RSVP-TE LSP between PE-2 and PE-4. The MPLS-TP LSP is protected using 1:1 linear protection, with a working path from PE-1 to PE-2, and a protect path from PE-1, through LSR P-3, to PE-2. The Ethernet PW for Epipe 10 connects an Ethernet SAP on port 1/1/5 on PE-1 to an Ethernet SAP on port 1/1/6 on PE-2. The PW for Epipe 20 connects an Ethernet SAP on port 1/1/6 on PE-1 to the VPLS on PE-4 and is switched between a static MPLS-TP segment and a dynamic targeted LDP (T-LDP) segment at PE-2. PE-2 thus acts as a gateway between the MPLS-TP domain and the IP/MPLS domain.

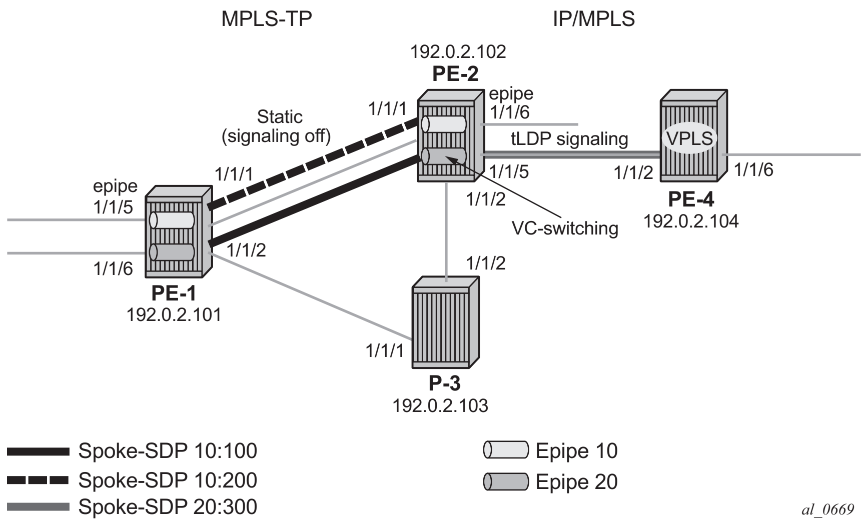

Figure 163: MPLS-TP Example Network Showing Services Detail

Figure 164shows the configuration process to be followed when setting up MPLS-TP.

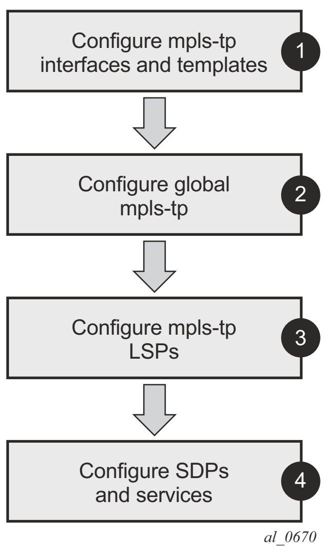

Figure 164: MPLS-TP Configuration Steps

Step 1. Configuration of MPLS-TP interfaces and templates.

MPLS-TP LSPs can use either numbered or unnumbered network IP interfaces, or unnumbered network interfaces that have been configured to operate without relying on IP routing. This non-IP interface type does not have an IP address associated with it and may be configured to have either a unicast, broadcast or multicast MAC address. The intent of using a broadcast or multicast MAC address is to enable a standard set of MAC addresses to be configured for a network without requiring any changes to the configuration of neighboring router interfaces each time an interface to which a router is connected is changed. Note that if a broadcast or multicast MAC address is used, then the operator should take care that only a point-to-point link is connected to the Ethernet port used by the interface. Otherwise, MPLS-TP packets may be replicated to each remote port to which the link is connected.

The non-IP network interface type is known as an unnumbered-mpls-tp interface. Only MPLS-TP can use this interface type. That is, other IP protocols are blocked from using it. Furthermore, ARP is not used for next hop resolution. This example uses unnumbered-mpls-tp interfaces.

Unnumbered MPLS-TP interfaces are configured on each network-facing interface for the nodes in the MPLS-TP domain, as shown below. This is done using the **unnumbered-mpls-tp** keyword at create time. In addition, the **static-arp unnumbered** command is used to set the next-hop unicast, broadcast or multicast MAC address of the interface. The system interface should also be configured. Numbered IP Network interfaces, bound to port 1/1/5 of PE-2 and port 1/1/2 of PE-4 are used for the IP/MPLS portion of the network in Figure 162.

```
A:PE-1>config>router
        interface "PE-1-P-3" unnumbered-mpls-tp
            port 1/1/2
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
        exit
        interface "PE-1-PE-2" unnumbered-mpls-tp
            port 1/1/1
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
        exit
        interface "system"
            address 192.0.2.101/32
            no shutdown
        exit
        autonomous-system 64511
        router-id 192.0.2.101


A:PE-2>config>router
        interface "PE-2-P-3" unnumbered-mpls-tp
            port 1/1/2
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
        exit
        interface "PE-2-PE-1" unnumbered-mpls-tp
            port 1/1/1
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
```

```
        exit
        interface "PE-2-PE-4"
            address 192.168.0.1/30
            port 1/1/5
            no shutdown
        exit
        interface "system"
            address 192.0.2.102/32
            no shutdown
        exit
        autonomous-system 64511
        router-id 192.0.2.102


A:P-3>config>router
        interface "system"
            address 192.0.2.103/32
            no shutdown
        exit
        interface "P-3-PE-1" unnumbered-mpls-tp
            port 1/1/1
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
        exit
        interface "P-3-PE-2" unnumbered-mpls-tp
            port 1/1/2
            static-arp unnumbered 01:00:5e:90:00:00
            no shutdown
        exit
        autonomous-system 64511
        router-id 192.0.2.103


A:PE-4>config>router
       interface "PE-4-PE-2"
            address 192.168.0.2/30
            port 1/1/2
            no shutdown
        exit
        interface "system"
            address 192.0.2.104/32
            no shutdown
        exit
        autonomous-system 6451165535
        router-id 192.0.2.104
```

Next, MPLS should be configured on each of the interfaces to be used by MPLS-TP. As an example, only PE-1 configuration is shown although a similar configuration is provisioned on PE-2 and P-3.

```
A:PE-1>config>router
        mpls
              mpls-tp
              exit
              interface "system"
```

```
    no shutdown
exit
interface "PE-1-PE-2"
    no shutdown
exit
interface "PE-1-P-3"
    no shutdown
exit
```

PE4 is an IP/MPLS only node so there is no MPLS TP configuration

```
A:PE-4>config>router
       mpls
            interface "system"
                no shutdown
            exit
            interface "PE-4-PE-2"
                no shutdown
            exit
```

Note that the mpls context must be in the **no shutdown** state to enable MPLS-TP.

Static labels are used by MPLS-TP LSPs and PWs. SR OS requires that a user reserves a range from the global label space for static labels. This prevents the labels being used by signaling protocols, such as RSVP. Static labels are reserved as shown by the following CLI command. As an example, the lower 100 labels (from 32, onwards) are reserved for static allocation to LSPs, and the next 100 are reserved for static allocation to PWs. This configuration should be repeated for every node that is an LER or LSR for MPLS-TP LSPs, although only the configuration for PE-1 is displayed.

```
A:PE-1>config>router
       mpls-labels
            static-labels max-lsp-labels 100 max-svc-labels 200
        exit
```

Next, one or more Bidirectional Forwarding Detection (BFD) templates are configured on the LERs. These templates are used to define BFD state machine parameters used for BFD Continuity Check (CC) on an LSP, including the transmit and receive timer intervals (in milliseconds). CPM network processor BFD is required if timer intervals as short as 10ms are used, but depending on the platform, 100ms BFD may use CPU based BFD (as shown in the example here).

```
config
   router
      bfd
         [no] bfd-template <name>
            [no] transmit-interval <transmit-interval>
            [no] receive-interval <receive-interval>
            [no] echo-receive <echo-interval>
            [no] multiplier <multiplier>
            [no] type <cpm-np>
            exit
```

A subset of these parameters is used by MPLS-TP BFD sessions, as follows:

transmit-interval *transmit-interval* and the **receive-interval** *receive-interval* — These are the transmit and receive timers for BFD packets. For MPLS-TP, these are the timers used by BFD CC packets. Values are in milliseconds: 10ms to 100,000ms, with 1ms granularity. Default 10ms for CPM3 or higher, 1 sec for other hardware. The minimum interval that can be supported is hardware dependent. For MPLS-TP BFD Connectivity Verification (CV) packets, a transmit interval of 1 sec is always used.

multiplier *multiplier* — Integer 3 – 20. Default: 3. The configured parameter is used for MPLS-TP CC BFD sessions. It is ignored for MPLS-TP combined CC/CV BFD sessions, and the default of 3 is used.

type cpm-np — This selects the CPM network processor as the local termination point for the BFD session. This is used by default for MPLS-TP. The CPM-NP type is needed to configure a transmit interval down to 10ms.

The following CLI illustrates the BFD template configuration at PE-1. Since default parameters are sufficient, only the bfd-template name is configured. Note that BFD templates use a begin/commit model for configuration. Create or modify a template with the **begin** statement. Changes to an existing template or the creation of a new template is not effected until the **commit** statement is entered.

```
A:PE-1>config>router
        bfd
            begin
            bfd-template "tp-bfd"
            exit
            commit
        exit
```

The following **info detail** command shows the values that are assigned by default.

```
A:PE-1>config>router>bfd# info detail
----------------------------------------------
            bfd-template "tp-bfd"
                no type
                transmit-interval 100
                receive-interval 100
                multiplier 3
                echo-receive 100
            exit
----------------------------------------------
```

Step 2. Configuration of Global MPLS-TP Parameters

MPLS-TP global parameters are configured under **config>router>mpls>mpls-tp**. These include the MPLS-TP identifiers for the node and the range of tunnel identifiers that should be reserved for MPLS-TP LSPs.

Node identifiers include the Global ID and the Node ID. The Node ID may be defined as an unsigned integer or use dotted quad notation (a.b.c.d), but the Node ID does not have to be a routable IP address.

The CLI tree for configuring the MPLS-TP identifiers for a node is as follows:

```
config
   router
     mpls
       mpls-tp
          [no] global-id <global-id>
          [no] node-id {<ipv4address> | <1...4,294,967,295>}
          [no] shutdown
          exit
```

The default value for the global-id is 0. This is used if the global-id is not configured. If an operator expects that inter-domain LSPs will be configured, then it is recommended to set the global ID to the local autonomous system number (ASN) of the node, as configured under **config>router**, to ensure that the combination of global-id and node-id is globally unique. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros.

The default value of the **node-id** is the system interface IPv4 address. The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value be configured.

In order to change the values, **config>router>mpls>mpls-tp** must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values are propagated to the system when a **no shutdown** is performed.

The following CLI shows the MPLS-TP node identifier configuration for PE-1. A similar configuration is implemented in all routers in this example, except that the node-ids must be different (PE-2 is 10.0.0.102 and P-3 is 10.0.0.103).

```
A:PE-1>config>router
         mpls
            mpls-tp
                global-id 64511
                node-id 10.0.0.101
```

Next, protection and OAM templates should be configured at the MPLS-TP LERs. A protection template defines the parameters of the linear protection state coordination mechanism. MPLS-TP Linear Protection is specified in RFC6378. It provides protection for an LSP using a working and a protect path. A protection state coordination (PSC) protocol is used by the LERs at each end of the protected LSP to coordinate whether the working or protect path is used for forwarding. BFD is run on both the working and protect paths.

The linear protection parameters include revertive or non-revertive behavior, the wait-to-restore timer, the rapid-psc-timer and the slow-psc-timer. The wait-to-restore timer (in seconds) defines the time to wait before reverting to the working path if, on restoration of connectivity, the revertive behavior is selected.

The following CLI tree is used to configure the protection template:

```
config
   router
      mpls
         mpls-tp
            protection-template <name>
              [no] revertive
              [no] wait-to-restore <interval>
              [no] rapid-psc-timer <interval>
              [no] slow-psc-timer <interval>
              exit
```

Refer to the CLI command descriptions in the MPLS User Guide for further details of these commands.

The OAM template defines generic proactive OAM parameters, such as BFD hold down and hold up timer values (which can be used to introduce some hysteresis if BFD bounces) and the BFD template to use.

The following CLI tree is used to configure the OAM template:

```
config
   router
      mpls
         mpls-tp
            [no] oam-template <name>
               [no] bfd-template <name>
               [no] hold-time-down <interval>
               [no] hold-time-up <interval>
               exit
```

Refer to the CLI command descriptions in the MPLS User Guide for further details of these commands.

MPLS-TP requires the reservation of a tunnel ID range, dedicated for the use of MPLS-TP LSPs. This range is reserved using the following CLI tree:

```
config
   router
      mpls
         mpls-tp
            [no] tp-tunnel-id-range <start-id> <end-id>
```

The default parameter values are used as shown below, where PE-1 and PE-2 have the same configuration:

```
A:PE-1>config>router mpls
            mpls-tp
                tp-tunnel-id-range 100 1000
                protection-template "tp-protect"
                exit
                oam-template "tp-oam"
                    bfd-template "tp-bfd"
                exit
                no shutdown
            exit
```

Step 3. Configuration of MPLS-TP LSPs

Once the global MPLS-TP parameters have been configured, the system is ready to configure MPLS-TP LSPs. An MPLS-TP LSP is configured under the **config>router>mpls>lsp** context.

Note that because LSP labels are statically configured, both ends of the LSP must be explicitly configured. The LSP paths must also be explicitly configured in the LSR nodes. MPLS-TP LSPs must use the **tp-lsp** and **source-tunnel-num** create time parameters.

The following commands are used to configure an MPLS-TP LSP at an LER:

```
config
   router
      mpls
         lsp <lsp-name> mpls-tp <src-tunnel-num>]
            to node-id {<a.b.c.d> | <1.. .4,294,967,295>}
            dest-global-id <global-id>
            dest-tunnel-number <tunnel-num>
            [no] working-tp-path
                lsp-num <lsp-num>
                in-label <in-label> [in-link <if-name>]
                out-label <out-label> out-link <if-name>
                          [next-hop <ipv4-address>]
                [no] mep
                   [no] oam-template <name>
                   [no] bfd-enable [cc | cc_cv]
                   [no] shutdown
                   exit
                [no] shutdown
                exit
             [no] protect-tp-path
                lsp-num <lsp-num>
                in-label <in-label> [in-link <if-name>]
                out-label <out-label> out-link <if-name>
                          [next-hop <ipv4-address> ]
                [no] mep
                   [no] protection-template <name>
                   [no] oam-template <name>
                   [no] bfd-enable [cc | cc_cv]
                   [no] shutdown
                   exit
                [no] shutdown
                exit
```

Refer to the CLI command descriptions in the MPLS User Guide for further details of these commands.

A working path and a protect path for LSP LSP-PE-1-P-3 must be configured between PE-1 and PE-2. Each LSP is configured with the full set of MPLS-TP identifiers required to build the LSP ID. Each working path and protect path must have an incoming label, outgoing label and outgoing link configured.

Each working path and protect path also includes a Maintenance Entity Group Endpoint (MEP) configuration, under which the applicable OAM template is configured. BFD is also enabled under the MEP context for the path. In this example, BFD operating in CC mode is enabled on the working and protect paths. Note that the Protection Template, containing parameters for linear protection, is only applied under the protect path context.

Figure 165 shows the LSP working and protect path label values configured at PE-1, PE-2 and P-3. Note that at each node the outgoing label must match the incoming label on the next hop for a given direction. At the LERs (PE-1 and PE-2), the incoming and outgoing label values for each LSP path are configured together. However, at the LSR (P-3), the label values for the label mapping between ingress and egress for each direction of the path (that is, forward and reverse) are configured together.

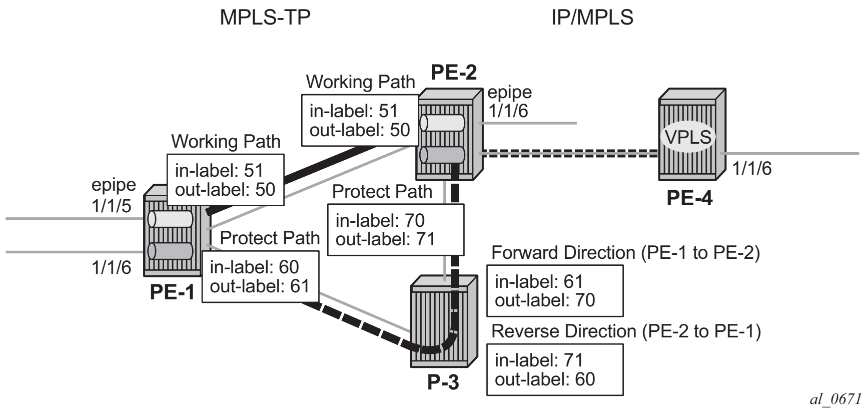

Figure 165: LSP Path Label Value Configurations

The following shows the LER LSP configuration of PE-1 and PE-2.

```
A:PE-1>config>router
        mpls
            lsp "LSP-PE-1-PE-2" mpls-tp 100
                to node-id 10.0.0.102
                dest-global-id 64511
                dest-tunnel-number 100
                working-tp-path
                    in-label 50
                    out-label 51 out-link "PE-1-PE-2"
                    mep
                        oam-template "tp-oam"
                        bfd-enable cc
                        no shutdown
                    exit
                    no shutdown
                exit
                protect-tp-path
                    in-label 60
                    out-label 61 out-link "PE-1-P-3"
                    mep
```

```
                        protection-template "tp-protect"
                        oam-template "tp-oam"
                        bfd-enable cc
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit


A:PE-2>config>router
        mpls
           lsp "LSP-PE-1-PE-2" mpls-tp 100
                to node-id 10.0.0.101
                dest-global-id 64511
                dest-tunnel-number 100
                working-tp-path
                    in-label 51
                    out-label 50 out-link "PE-2-PE-1"
                    mep
                        oam-template "tp-oam"
                        bfd-enable cc
                        no shutdown
                    exit
                    no shutdown
                exit
                protect-tp-path
                    in-label 70
                    out-label 71 out-link "PE-2-P-3"
                    mep
                        protection-template "tp-protect"
                        oam-template "tp-oam"
                        bfd-enable cc
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
```

Since this example requires a protect path to be switched via P-3, a transit path must be configured in P-3. The CLI tree for configuring MPLS-TP transit paths is as follows:

```
config
   router
      mpls
         mpls-tp
            transit-path <path-name>
               [no] path-id {lsp-num <lsp-num>|working-path|protect-path

                  [src-global-id <global-id>]
                  src-node-id {<ipv4address> | <1.. .4,294,967,295>}
```

```
            src-tunnel-num <tunnel-num>
            [dest-global-id <global-id>]
            dest-node-id {<ipv4address> | <1.. .4,294,967,295>}
            [dest-tunnel-num <tunnel-num>]}

         forward-path
            in-label <in-label> out-label <out-label>
            out-link <if-name> [next-hop <ipv4-next-hop>]
            exit
         reverse-path
            in-label <in-label> out-label <out-label>
            [out-link <if-name> [next-hop <ipv4-next-hop>]]
             exit
         [no] shutdown
```

Refer to the CLI command descriptions in the MPLS User Guide for further details of these commands.

The CLI configuration for the forward and reverse directions of the transit path (that is, the protect path of the LSP) at P-3 is as follows:

```
A:P-3>config>router
   mpls
      mpls-tp
         transit-path "LSP-PE-1-PE-2"
            forward-path
               in-label 61 out-label 70 out-link "P-3-PE-2"
               exit
            reverse-path
               in-label 71 out-label 60 out-link "P-3-PE-1"
                exit
            path-id src-global-id 64511 src-node-id 10.0.0.101 src-tunnel-num 100 dest-
global-id 64511 dest-node-id 10.0.0.102 dest-tunnel-num 100 lsp-num 2
            no shutdown
            exit
       no shutdown
       exit
    exit
```

The example also requires an LSP across the IP/MPLS network to backhaul traffic from PE-2 at the edge of the MPLS-TP network to the VPLS service hosted in PE-4. An RSVP LSP is configured at PE-2 for this purpose, as follows:

```
A:PE-2>config>router
        mpls
            path "completely-loose-path"
                no shutdown
            exit
            lsp "LSP-PE-2-PE-4"
                to 192.0.2.104
                primary "completely-loose-path"
                exit
                no shutdown
            exit
```

At this point in the configuration process, it is recommended to check the MPLS-TP LSP configuration and operation of BFD and linear protection.

First, check that the BFD sessions on both the working and protect paths are up:

```
*A:PE-1# show router bfd session
===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Interface/Lsp Name           State                Tx Intvl  Rx Intvl  Multipl
  Remote Address/Info        Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                     LAG ID
-------------------------------------------------------------------------------
wp::LSP-PE-1-PE-2            Up (3)               100       100       3
    64511::10.0.0.102        mplsTp               684       692       central
pp::LSP-PE-1-PE-2            Up (3)               100       100       3
    64511::10.0.0.102        mplsTp               164       167       central
-------------------------------------------------------------------------------
No. of BFD sessions: 2
===============================================================================
===============================================================================
```

Next, check the currently active path. This can be done using the oam lsp-trace command. Note that the static option must be specified for MPLS-TP LSPs.

```
*A:PE-1# oam lsp-trace static "LSP-PE-1-PE-2"
lsp-trace to LSP-PE-1-PE-2: 0 hops min, 0 hops max, 100 byte packets
1  GlobalId 64511 NodeId 10.0.0.102
   rtt=0.648ms rc=3(EgressRtr)
```

This shows that data packets currently follow the working path of the LSP (no transit node is shown).

In order to test the operation of linear protection, the port used by the working path can be shutdown, and the BFD session state checked again:

```
*A:PE-1# configure port 1/1/1 shutdown
*A:PE-1# show router bfd session
===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Interface/Lsp Name           State                Tx Intvl  Rx Intvl  Multipl
  Remote Address/Info        Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                     LAG ID
-------------------------------------------------------------------------------
wp::LSP-PE-1-PE-2            Down (1)             1000      100       3
    64511::10.0.0.102        mplsTp               3171      3170      central
pp::LSP-PE-1-PE-2            Up (3)               100       100       3
    64511::10.0.0.102        mplsTp               2727      2730      central
-------------------------------------------------------------------------------
```

```
No. of BFD sessions: 2
===============================================================================
```

Execute LSP trace again to check that the LSP has failed over to use the protect path:

```
*A:PE-1# oam lsp-trace static "LSP-PE-1-PE-2"
lsp-trace to LSP-PE-1-PE-2: 0 hops min, 0 hops max, 100 byte packets
1  GlobalId 64511 NodeId 10.0.0.103
   rtt=0.868ms rc=8(DSRtrMatchLabel)
2  GlobalId 64511 NodeId 10.0.0.102
   rtt=1.15ms rc=3(EgressRtr)
```

This shows that packets are now forwarded via the protect path through P-3, which has Node ID 10.0.0.103.

Finally bring the LSP back to the working path by bringing port 1/1/1 up, and either waiting for the LSP to revert to the working path or forcing it onto the working path and clearing the revert timer by executing a tools command as follows:

```
*A:PE-1# tools perform router mpls tp-tunnel force "LSP-PE-1-PE-2"
*A:PE-1# tools perform router mpls tp-tunnel clear "LSP-PE-1-PE-2"
*A:PE-1# oam lsp-trace static "LSP-PE-1-PE-2"
lsp-trace to LSP-PE-1-PE-2: 0 hops min, 0 hops max, 100 byte packets
1  GlobalId 64511 NodeId 10.0.0.102
   rtt=0.637ms rc=3(EgressRtr)
```

Step 4. Step 4: Configuration of SDPs and Services

Services can be configured to use MPLS-TP LSPs once the LSP configuration is completed. SDPs and services are configured in a similar manner to those using static-labelled pseudowires without MPLS-TP.

Distributed services are configured to use MPLS-TP with the following steps:

- Configure an SDP with signaling off. With signaling off, the SDP far-end may then be configured as an MPLS-TP node-id or an IPv4 address. SDP keep-alive should be disabled.
- Configure the service, including the spoke-sdp using the SDP. To use MPLS-TP, the spoke-sdp must have statically assigned ingress and egress labels, the control-word must be enabled, and it must have an MPLS-TP identifier for the PW (the PW Path ID) configured. This is comprised of two parts, a Source Attachment Individual Identifier (SAII) and a Target Attachment Individual Identifier(TAII), both of which must be configured. Control channel status signaling may also be configured to support PW status signaling on the static MPLS-TP PW.

In this example, an SDP is configured to use the MPLS-TP LSP from PE-1 to PE-2, which will act as a transport for the static MPLS-TP PWs corresponding to Epipe 10 and Epipe 20. A further SDP is configured for the targeted LDP (T-LDP) PW segment corresponding to Epipe 20 between PE-2 and PE-4.

Note that Epipe 10 belongs to customer 1, and Epipe 20 belongs to customer 2 in this example.

The following CLI shows the SDP between PE-1 and PE-2 and the SDP between PE-2 and PE-4:

```
A:PE-1>config
    service
        sdp 10 mpls create
            signaling off
            far-end node-id 10.0.0.102 global-id 64511
            lsp "LSP-PE-1-PE-2"
            keep-alive
                shutdown
            exit
            no shutdown
        exit

A:PE-2>config
   service
        sdp 10 mpls create
            signaling off
            far-end node-id 10.0.0.101 global-id 64511
            lsp "LSP-PE-1-PE-2"
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 20 mpls create
```

```
            far-end 192.0.2.104
            lsp "LSP-PE-2-PE-4"
            keep-alive
                no shutdown
            exit
            no shutdown
        exit

A:PE-4>config
   service
       sdp 20 mpls create
            far-end 192.0.2.102
            lsp "LSP-PE-2-PE-4"
            keep-alive
                no shutdown
            exit
            no shutdown
        exit
```

Next, configure the services that will use the MPLS-TP LSPs.

The service configuration CLI tree for an Epipe service using MPLS-TP is as follows:

```
config
   service
      epipe
         [no] spoke-sdp sdp-id[:vc-id]
              [no] hash-label
              [no] standby-signaling-slave
         [no] spoke-sdp sdp-id[:vc-id] [vc-type {ether|vlan}]
             [create] [vc-switching] [no-endpoint | {endpoint [icb]}]
            egress
               vc-label <out-label>
            ingress
               vc-label <in-label>
            [no] control-word
            [no] pw-path-id
               agi <agi>
               saii-type2 <global-id:node-id:ac-id>
               taii-type2 <global-id:node-id:ac-id>
               exit
            control-channel-status
               [no] acknowledgment
               [no] refresh-timer <value>
               [no] request-timer <value> retry-timer <value> [timeout-multiplier <value>]
               [no] shutdown
               exit
```

Refer to the CLI command descriptions in the user guides for further details of these commands.

The following CLI examples show the Epipe service configuration at PE-1, PE-2, and the VPLS spoke-sdp termination point at PE-4.

```
A:PE-1>config
     service
        epipe 10 customer 1 create
            sap 1/1/5 create
            exit
            spoke-sdp 10:100 create
                ingress
                    vc-label 150
                exit
                egress
                    vc-label 151
                exit
                control-word
                pw-path-id
                    saii-type2 64511:10.0.0.101:1
                    taii-type2 64511:10.0.0.102:1
                exit
                control-channel-status
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        epipe 20 customer 2 create
            sap 1/1/6 create
            exit
            spoke-sdp 10:200 create
                ingress
                    vc-label 200
                exit
                egress
                    vc-label 201
                exit
                control-word
                pw-path-id
                    saii-type2 64511:10.0.0.101:2
                    taii-type2 64511:10.0.0.102:2
                exit
                control-channel-status
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
```

At PE-2, Epipe 10 terminates on a SAP on port 1/1/6, while Epipe 20 is switched between a static MPLS-TP PW segment (spoke-sdp 10:200) and a T-LDP signaled PW segment (spoke-sdp 20:300) for backhaul to the remote PE-4 containing the VPLS service.

At PE-4, the T-LDP signaled PW segment for Epipe 20 is terminated on a VPLS service:

```
A:PE-4>config
    service
        vpls 1 customer 2 create
            stp
                shutdown
            exit
            sap 1/1/6 create
            exit
            spoke-sdp 20:300 create
                control-word
                no shutdown
            exit
            no shutdown
        exit
```

Epipe 10 uses a static MPLS-TP PW from end to end, which can be tested using the vccv-ping command at PE-1, as follows:

```
A:PE-1#  oam vccv-ping static 10:100
VCCV-PING 10:100 84 bytes MPLS payload
Seq=1, send from intf PE-1-PE-2
       send from lsp LSP-PE-1-PE-2
       reply via Control Channel
       src id tlv received: GlobalId 64511 NodeId 10.0.0.102
       cv-data-len=44 rtt=0.992ms rc=3 (EgressRtr)

---- VCCV PING 10:100 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 0.992ms, avg = 0.992ms, max = 0.992ms, stddev = 0.000ms
```

The operation of control channel status signaling can also be tested for this Epipe, as follows:

Shutdown the port that the SAP on PE-2 is using:

```
*A:PE-2# configure port 1/1/6 shutdown
```

The PW peer status bits for the spoke-sdp for Epipe 10, signaled using control channel status signaling, can be displayed at node PE-1 using the following command (note that some of the show command output has been removed for brevity). The peer PW status bits are shown in **bold** in the output below.

```
A:PE-1# show service id 10 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 07/05/2014 00:01:52
```

```
Last Mgmt Change  : 07/05/2014 00:01:09
Admin State       : Up                   Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                    SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

<snipped>
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 10:100  -(10.0.0.102: 64511)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 10:100                     Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : Ether                      VC Tag            : n/a
Admin Path MTU    : 0                          Oper Path MTU     : 8914
Delivery          : MPLS
Far End           : 10.0.0.102: 64511
Tunnel Far End    : n/a                        LSP Types         : MPLSTP
Hash Label        : Disabled                   Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled

Admin State       : Up                         Oper State        : Up
Acct. Pol         : None                       Collect Stats     : Disabled
Ingress Label     : 150                        Egress Label      : 151
Ingr Mac Fltr-Id  : n/a                        Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id   : n/a                        Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id : n/a                        Egr IPv6 Fltr-Id  : n/a
Admin ControlWord : Preferred                  Oper ControlWord  : True
Admin BW(Kbps)    : 0                          Oper BW(Kbps)     : 0
Last Status Change : 07/05/2014 00:01:52      Signaling         : None
Last Mgmt Change  : 07/05/2014 00:01:09       Force Vlan-Vc     : Disabled
Endpoint          : N/A                        Precedence        : 4
PW Status Sig     : Enabled
Class Fwding State : Down
Flags             : None
Local Pw Bits     : None
Peer Pw Bits      : lacIngressFault lacEgressFault
Peer Fault Ip     : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
...
```

Epipe 20 uses a static MPLS-TP PW from PE-1 to PE-2, identified by a static PW Forwarding Equivalence Class (FEC), and a T-LDP segment with FEC128 from PE-2 to PE-4. Therefore the target FEC used for a vccv-ping command from PE-1 to PE-4 is different from the local FEC for the PW at PE-1. VCCV-trace provides a useful tool to test the resulting multi-segment PW (MS-PW). Note that the same associated channel type must be used for both segments. This is the IPv4 channel.

```
A:PE-1# oam vccv-trace static 10:200 assoc-channel ipv4 detail
VCCV-TRACE 10:200  with 116 bytes of MPLS payload
```

```
1  192.0.2.102 GlobalId 64511 NodeId 10.0.0.102
   rtt=0.733ms rc=8(DSRtrMatchLabel)
   Next segment: VcId=300 VcType=Ether Source=192.0.2.102 Remote=192.0.2.104
2  192.0.2.104  rtt=1.90ms rc=3(EgressRtr)
```

The system supports the interworking of control channel status on a static MPLS-TP PW segment with T-LDP-signaled PW status on a T-LDP PW segment. This can be tested as follows.

Shutdown the port that the spoke sdp on PE-4 is using:

```
A:PE-4# configure port 1/1/2 shutdown
```

The PW peer status bits for the spoke-sdp for Epipe 20 can then be displayed at node PE-1 using the following command (note that some of the show command output has been removed for brevity). The peer PW status bits are shown in **bold** in the output below.

```
A:PE-1# show service id 10 all
A:PE-1>show>service>id# all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 20                  Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 2                   Creation Origin   : manual
Last Status Change: 07/05/2014 00:01:52
Last Mgmt Change  : 07/05/2014 00:01:09
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

<snipped>
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 10:200  -(10.0.0.102: 64511)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id           : 10:200                     Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type          : Ether                      VC Tag            : n/a
Admin Path MTU   : 0                          Oper Path MTU     : 8914
Delivery         : MPLS
Far End          : 10.0.0.102: 64511
Tunnel Far End   : n/a                        LSP Types         : MPLSTP
Hash Label       : Disabled                   Hash Lbl Sig Cap  : Disabled
Oper Hash Label  : Disabled

Admin State      : Up                         Oper State        : Up
Acct. Pol        : None                       Collect Stats     : Disabled
```

```
Ingress Label      : 200                    Egress Label      : 201
Ingr Mac Fltr-Id   : n/a                    Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                    Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                    Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Preferred              Oper ControlWord  : True
Admin BW(Kbps)     : 0                      Oper BW(Kbps)     : 0
Last Status Change : 07/05/2014 00:01:52    Signaling         : None
Last Mgmt Change   : 07/05/2014 00:01:09    Force Vlan-Vc     : Disabled
Endpoint           : N/A                    Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : psnIngressFault psnEgressFault
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Peer Vccv CC Bits  : None

...
```

Conclusion

Release 11.0.R4 of SR OS introduced extensive MPLS Transport Profile (MPLS-TP) capabilities. MPLS-TP is intended to allow MPLS to be operated in a similar manner to existing transport technologies, with in-band proactive and on-demand operations and maintenance (OAM), and protection mechanisms that do not rely on a control plane to operate. The 7x50 can operate both as an LER and LSR for MPLS-TP LSPs, and as a T-PE and S-PE for PWs with MPLS-TP OAM. The 7x50 can therefore act as a node within an MPLS-TP network, or as a gateway between MPLS-TP and IP/MPLS domains.

This example has illustrated a simple configuration, demonstrating the role of the 7x50 as an LER and LSR for MPLS-TP LSPs, and how its already extensive multi-service capabilities can be extended over an MPLS-TP network and between MPLS-TP and IP/MPLS networks.