



CHAPTER 8

Managing MPLS Traffic Engineering Services

This chapter contains a detailed description of the Cisco Prime Provisioning Traffic Engineering Management (TEM) product, including the various features, the GUI, and the step-by-step processes needed to perform various traffic engineering management tasks.

This chapter includes the following sections:

- [Getting Started, page 8-1](#)
- [TE Network Discovery, page 8-10](#)
- [TE Resource Management, page 8-20](#)
- [Basic Tunnel Management, page 8-27](#)
- [Advanced Primary Tunnel Management, page 8-44](#)
- [Protection Planning, page 8-58](#)
- [TE Traffic Admission, page 8-67](#)
- [Administration, page 8-70](#)
- [TE Topology, page 8-80](#)
- [Sample Configlets, page 8-87](#)
- [Warnings and Violations, page 8-98](#)
- [Document Type Definition \(DTD\) File, page 8-108](#)
- [Traffic Engineering Management Concepts, page 8-111](#)

Getting Started

This section describes the installation procedure for Prime Provisioning. The general installation procedure for Cisco Prime Provisioning (Prime Provisioning) is described in the [Cisco Prime Provisioning 6.4 Installation Guide](#).

It includes the following sections:

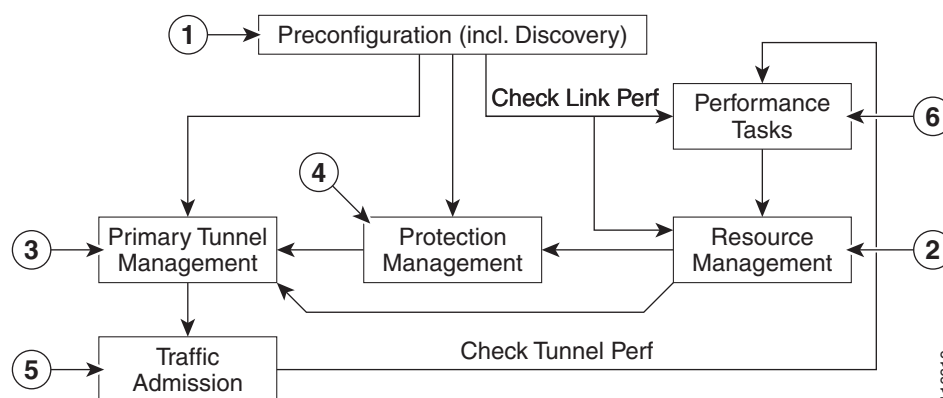
- [Prerequisites and Limitations, page 8-3](#)
 - [General Limitations, page 8-3](#)
 - [Feature-Specific Prerequisites and Limitations, page 8-3](#)
 - [Non-Cisco Devices and TEM, page 8-4](#)
 - [Supported Platforms, page 8-4](#)

- [Error Messages](#), page 8-4
- [Preconfiguration Process Overview](#), page 8-4
- [TEM Setup and Installation](#), page 8-6
 - [Editing DCPL Properties \(Optional\)](#), page 8-7
- [Creating a TE Provider](#), page 8-7

Process Overview

The main components and flows in TEM are shown in [Figure 8-1](#).

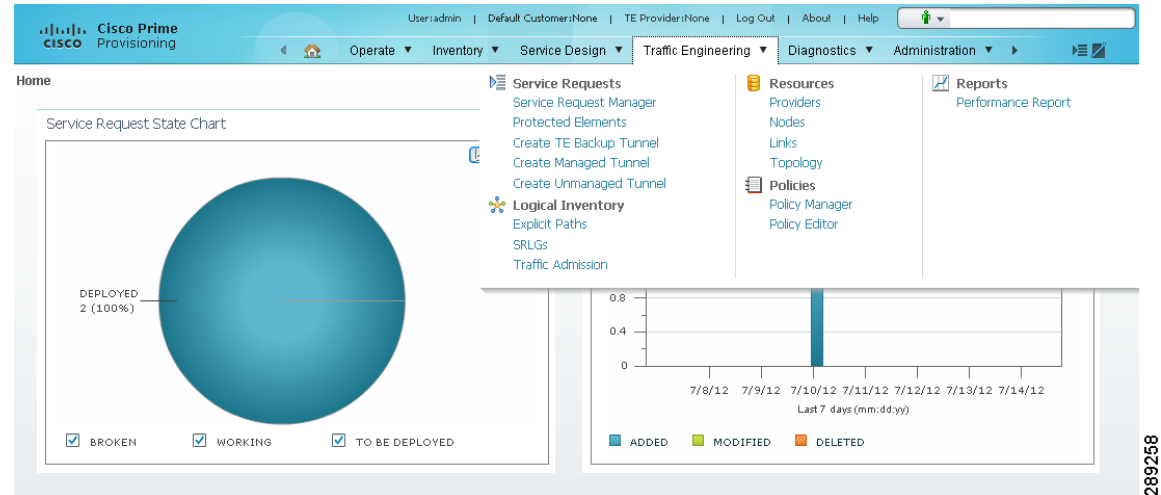
Figure 8-1 Main Process Flows in TEM



The illustration includes the following components:

1. **Preconfiguration**—Sets up key parameters that enable the system to collect TE network information (TE Discovery) and subsequently deploy TE configurations on the chosen network. (See [Getting Started](#), page 8-1)
2. **Resource Management**—Tuning of certain properties on the TE interfaces to optimize the tunnel placement. (See [TE Resource Management](#), page 8-20)
3. **Primary Tunnel Management**—Create and manage primary tunnels, either unmanaged (see [Basic Tunnel Management](#), page 8-27) or managed. (see [Basic Tunnel Management](#), page 8-27 or [Advanced Primary Tunnel Management](#), page 8-44)
4. **Protection Management**—Protect selected elements in the network (links, routers, or SRLGs) against failure. (See [Advanced Primary Tunnel Management](#), page 8-44)
5. **Traffic admission**—Assign traffic to traffic-engineered tunnels. (See [TE Traffic Admission](#), page 8-67)
6. **Performance Tasks**—Calculates interface/tunnel bandwidth utilization using the Simple Network Management Protocol (SNMP). (See [Administration](#), page 8-70)

The Traffic Engineering menu options in the Prime Provisioning user interface are shown in [Figure 8-2](#).

Figure 8-2 Traffic Engineering Menu Options

289258

Prerequisites and Limitations

The current release of Prime Provisioning involves certain prerequisites and limitations, which are described in this section.

See the [Cisco Prime Provisioning 6.4 Installation Guide](#) for general system recommendations.

General Limitations

The present release of Prime Provisioning has the following limitations:

- Although concurrent use of Prime Provisioning is supported in the Planning portion of the current implementation (see the section [Multiple Concurrent Users](#), page 8-114), multiple browsers on the same machine are still not recommended due to a limitation in Browser Session Attributes.
- JRE version 1.6.0_07 or higher should be installed on the client computer for launching Java applications and Applets. This can be done via Java's Control Panel. If you do not already have Java installed, you can use the links on the Topology Tool page to install the version that is bundled with Prime Provisioning.
- If your repository predates the ISC 4.1 release and has been upgraded to a 4.1 or later repository, you need to run a TE Discovery task to collect software version information from the devices before deploying service requests.
- Let issued service requests finish deployment before issuing other service requests to avoid conflicts. This is described in more detail in the tunnel provisioning sections.

Feature-Specific Prerequisites and Limitations

Prime Provisioning has the following feature-specific prerequisites and limitations:

- Some features might only be available with a particular license. In addition, the number of nodes provided by the license limits the size of the network. For more information, see [Traffic Engineering Management Concepts](#), page 8-111.

- A number of specific requirements are associated with the TE Discovery task. These are described in [TE Discovery Prerequisites and Limitations, page 8-12](#).
- Prime Provisioning manages a single OSPF area or IS-IS level. Prime Provisioning also supports multiple OSPF areas, however it does not discover tunnels between areas. Each OSPF area is mapped to a TE provider and is discovered area by area independently.
- Prime Provisioning only supports MPLS-TE topology with point-to-point links.

Non-Cisco Devices and TEM

Prime Provisioning does not manage non-Cisco devices and Prime Provisioning cannot be used to provision them.

Prime Provisioning will, however, discover non-Cisco devices and store them in the repository. Tunnels can be run through these devices, the bandwidth consumed can be accounted for, but the devices are not otherwise managed by Prime Provisioning. TE tunnels originating from non-Cisco devices will not be discovered.

Sorting can be performed on different attributes in various parts of the Prime Provisioning GUI. However, due to the added support for non-Cisco devices, sorting cannot be performed on Device Name and MPLS TE ID in the TE Nodes List window.

Supported Platforms

For supported devices and IOS platforms, see the [Cisco Prime Provisioning 6.4 Installation Guide](#).

Error Messages

Warnings and violations that are invoked when using the TE planning tools in Prime Provisioning are documented in [Warnings and Violations, page 8-98](#)

Elixir warning messages might appear when performing deployments in Prime Provisioning:

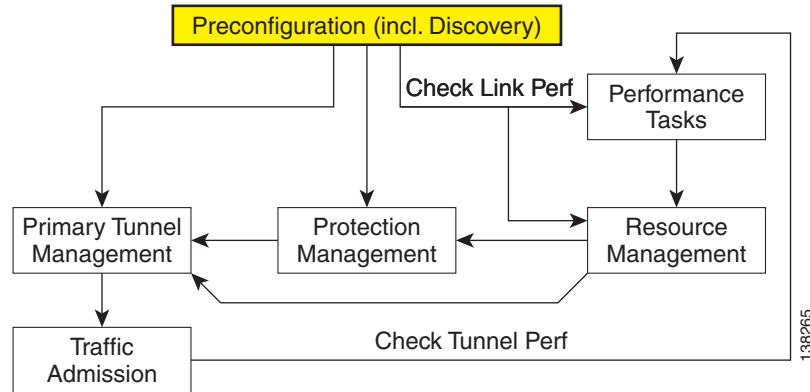
```
WARNING Elixir.ServiceBlade Unable to load support matrix for the platform or platform family. The default support matrix is loaded instead for role: TunnelHead.  
WARNING Elixir.ConfigManager Attribute - lockdown of Command - Tunnel_PathOption can NOT be retrieved from the input SR - SKIPPING.
```

The deployments will, however, be successful and these messages can be safely ignored.

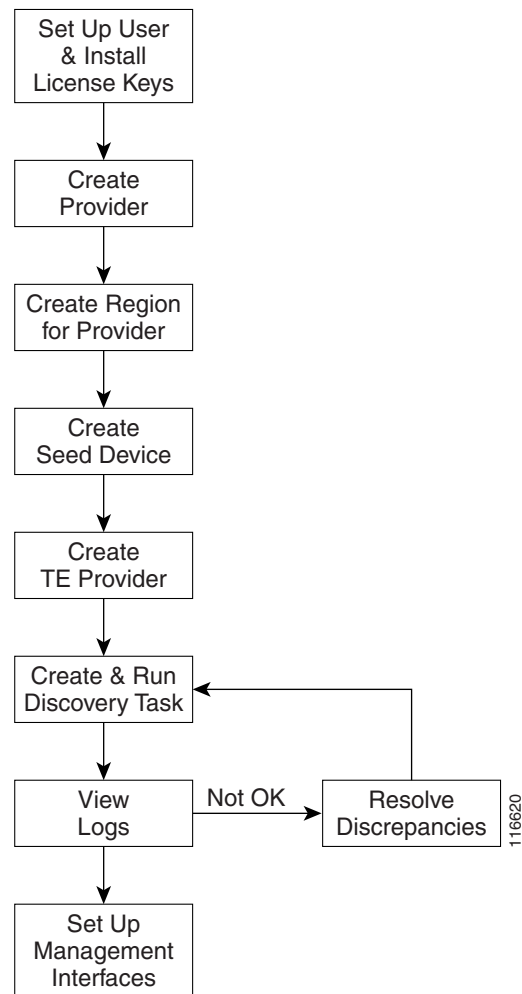
Preconfiguration Process Overview

The preconfiguration process sets up key parameters that enable the system to collect TE network information and subsequently deploy TE configurations on the chosen network.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning the preconfiguration steps take place.

Figure 8-3 Prime Provisioning Process Diagram - Preconfiguration

The different steps in the preconfiguration process are provided in [Figure 8-4](#).

Figure 8-4 Preconfiguration Process

Before commencing the preconfiguration process, MPLS-TE needs to be enabled on the network devices by making sure that the IP addresses used as devices' TE IDs are accessible from the management station (this step is not supported by TEM).

The preconfiguration process includes the following steps:

1. **Set up new user and install license keys**—To run the TEM blade of Prime Provisioning, it is necessary to create a new user and install license keys. These keys allow you to view and manage the TE tunnels and resources using Prime Provisioning. (See [TEM Setup and Installation](#), page 8-6)
2. **Create a provider**—The provider is a concept designed to allow many different operators to work on Prime Provisioning simultaneously, each working on different networks. Thus, each provider has to be defined and used as a reference operator for future work on the system. (To create a provider, see [Providers](#), page 2-15.)
3. **Create a region for the provider**—The region is important because a single provider could have multiple networks. The region is used as a further level of differentiation to allow for such circumstances. (To create a region, see [Provider Regions](#), page 2-16.)
4. **Create a seed device**—This IOS or IOS XR device will be the seed router for TE Discovery. The network discovery process uses the seed router as an initial communication point to discover the MPLS TE network topology. (To create a seed router, see [Devices](#), page 2-1.)
5. **Create a TE Provider**—Providers can be defined as TE provider, if they are supporting MPLS TE in their network. To enable a TE network to be managed, it is necessary to create a TE provider. All TE related data associated with a given network is stored under a unique TE provider. A provider and region uniquely define a TE provider (See [Creating a TE Provider](#), page 8-7.)
6. **Run a TE Discovery Task**—Discover the TE network for a particular TE provider to populate the repository with a view to creating primary and backup tunnels. (See [TE Network Discovery](#), page 8-10.)
7. **Set Up Management Interfaces**—Set up management interfaces for discovered devices or update server host file with resolution for all discovered devices. This step is only necessary if the devices in the TE network are not accessible via their hostnames (See [Setting Up Management Interfaces](#), page 8-19.)

**Note**

If Telnet is selected to communicate with the seed router, Telnet must also be used for the other network devices. Likewise, if SSH is selected for the seed router, SSH must be used for all other devices.

TEM Setup and Installation

Before setting up Prime Provisioning, the Prime Provisioning software must be installed. To do so, see the [Cisco Prime Provisioning 6.4 Installation Guide](#).

To set up a new Prime Provisioning user, one or more users with a TE role must be created. For step by step instructions, see [Cisco Prime Provisioning 6.4 Administration Guide](#).

Licensing information, including the Prime Provisioning licensing options and the procedure needed to install licenses is described in [Cisco Prime Provisioning 6.4 Administration Guide](#).

Editing DCPL Properties (Optional)

The Prime Provisioning Dynamic Component Properties Library (DCPL) includes a wide variety of properties that are accessible from the GUI, some of which can be modified.



Warning

Do not attempt to modify the DCPL properties unless you fully understand the implications.

In the Prime Provisioning GUI, the DCPL properties are found in **Administration > Hosts**. Check a check box for a specific host and click the **Config** button.

The DCPL properties pertaining to TEM are found in the following folders:

- **Provisioning > Service > TE**
- **TE**
- **TE Topology**

Creating a TE Provider

Before TE Discovery or any manipulation of TE data can take place, at least one TE provider has to be created. For example, an OSPF area can be assigned as a TE provider. Prior to this, a provider and a region for that provider must have been set up (see [Preconfiguration Process Overview, page 8-4](#)).

One region can be assigned as the default region as a place for discovered routers. These routers can then subsequently be placed in any region. For more information, see the section multiple hosts in [Cisco Prime Provisioning 6.4 Administration Guide](#).

To create a TE provider, use the following steps:

Step 1 Choose **Traffic Engineering > Providers**.

The TE Providers window appears.

Step 2 Click **Create** to create a TE provider.

The Create/Edit TE Provider window in [Figure 8-5](#) appears.

Figure 8-5 Create/Edit TE Provider

Create/Edit TE Provider

TE Provider Info:	
TE Provider *	te_provider2
Provider *	Select Provider1
TE Provider Area:	
TE Area	100
Primary Route Generation Parameters:	
Default Primary RG Timeout (sec) *	100
Backup Route Generation Parameters:	
Backup RG Timeout (sec) *	1000
FRR Protection Type *	<input checked="" type="radio"/> Sub Pool <input type="radio"/> Any Pool
Default Link Speed Factor *	1.00
Minimum Bandwidth Limit (Kbps) *	10
Max. Load Balancing Tunnel Count *	1
Discovery Default Parameters:	
Default Region for TE Devices *	Select Region4
Customer for Primary Tunnels:	Select
Select as default TE provider:	<input type="checkbox"/>
<div>Save Cancel</div>	

Note: * - Required Field

The Create/Edit TE Provider window includes the following fields:

- **TE Area**—OSPF area assigned to the TE provider. This can be any positive integer from 0 to 4294967295 or a dot notation address of the form x.x.x.x where x is a number between 0 and 255.
- **Default Primary RG Timeout**—Default computation timeout for primary tunnels.
- **Backup RG Timeout**—Computation timeout per element for backup tunnels (for each protected element, the timer is reset to zero before the Prime Provisioning attempts to protect it).
- **FRR Protection Type**—Fast Re-Route (FRR) protection type:
 - **Sub Pool**—Protect only sub pool primary tunnels.
 - **Any Pool**—Protect both sub pool and global pool primary tunnels.

For a definition of pool types, see the section on bandwidth pools in [Traffic Engineering Management Concepts, page 8-111](#).

- **Default Link Speed Factor**—Default multiplication factor to be applied to the link speed in order to determine move affected tunnels. that needs to be protected. The link's bandwidth is multiplied by the link speed factor, then the RSVP bandwidth reserved for the link (sub pool or global pool depending on the FRR protection type) is subtracted, and the resulting bandwidth is then available to FRR backup tunnels.

Interpretation of the link speed factor:

> 1.0 (overbooking)—more backup bandwidth than the link has available.

< 1.0 (underbooking)—less backup bandwidth than the link has available.

- **Minimum Bandwidth Limit**—Minimum bandwidth allowed for backup tunnels.
- **Max. Load Balancing Tunnel Count**—This is the maximum number of backup tunnels needed to protect a flow through a protected element. Here, a flow is defined as follows:

There are two flows in a protected link, one in each of the directions that traffic can flow. For a node, the number of flows depends on the number of neighbouring nodes for a particular node. There is a flow for each neighbour pair. So a node with 3 neighbours, A, B, and C, has 6 flows through it – A->B, A->C, B->A, B->C, C->A, C->B.

- **Default Region for TE Devices**—The default provider region is the one assigned by TE Discovery to a newly discovered device. If the device already exists in the repository and has a region defined, TE Discovery keeps that setting. It is possible to change the region of a device after TE Discovery.
- **Customer for Primary Tunnels**—Name of customer for primary TE tunnels.

Step 3 In the **TE Provider** field, enter a name for the new TE provider.

Step 4 To select a provider to be this TE provider, click the **Select** button next to the **Provider** field.

The Select Provider window appears.

Step 5 Select the desired provider using the radio buttons or search for a provider with search criteria matching a provider name and click **Find**.

Step 6 Click **Select** to select the desired provider.

The Select Provider window closes. The selected provider name is displayed in the **Provider** field.

Step 7 In the **TE Area** field, specify the number of the OSPF area to act as TE area.

Both dot notation and decimal notation are supported for the area identifier.



Note The **TE Area** field can be left blank if the seed router used for TE Discovery is not an Area Border Router, and it will be automatically populated on discovery.

Depending on the seed router used for TE Discovery, the area identifier should be set as follows:

- **Seed router is an ABR:** The area identifier field in TE provider must be set to indicate which of the two or more areas on the ABR is to be discovered.
- **Seed router is NOT an ABR:** Leave blank.



Note If you do not set the Area Identifier in TE Provider, TE Discovery will set it. After it is set, it cannot be changed.

Step 8 Add primary and backup route generation parameters.

When the FRR (Fast Re-Route) protection type is equal to Sub Pool, the backup tunnels generated by the tool will protect only the sub pool primary tunnels. When it is equal to Any Pool, the backup tunnels generated by the tool will protect both sub pool and global pool primary tunnels.

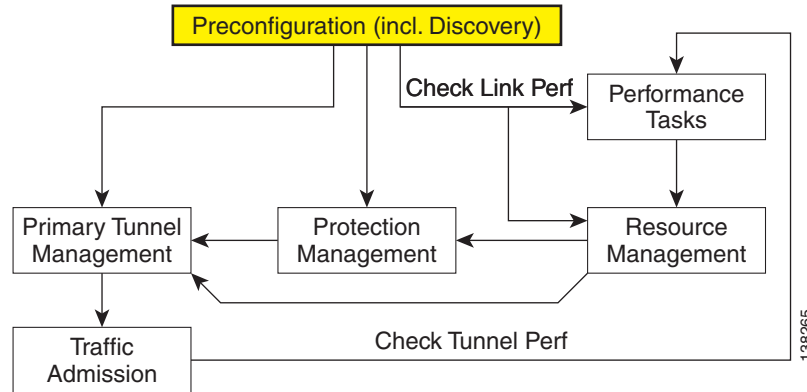
For more information on Fast Re-Route (FRR) protection pools, see the section on bandwidth pools in [Traffic Engineering Management Concepts, page 8-111](#).

- Step 9** Fill in the remaining required fields (marked ‘*’) and any optional fields as desired.
- Step 10** For the required **Default Region for TE Devices** field, click the corresponding **Select** button.
The Region for Create TE Provider window appears.
- Step 11** Select the desired region using the radio buttons.
- Step 12** Click **Select** to select the desired default region.
The Region for Create TE Provider window closes. The selected region name is displayed in the **Default Region for TE Devices** field.
- Step 13** For the optional **Customer for Primary Tunnels** field, click the corresponding **Select** button.
The Customer for Create TE Provider window appears.
- Step 14** If desired, select a customer using the radio buttons or search for a customer by entering customer search criteria in the **Show Customers with Customer Name matching** field and click **Find**.
- Step 15** Click **Select** to select the desired customer.
The Select Customer for Create TE Provider window closes. The selected customer name is displayed in the **Customer for Primary Tunnels** field of the Create/Edit TE Provider window.
- Step 16** Click **Save**.
The created TE provider appears in the TE Provider window and can now be used to perform TE discovery and other TE functions.
- To switch between TE providers, go to the top of the Prime Provisioning window above the menu toolbar ([Figure 8-2](#)) and click the **TE Provider** link.
-

TE Network Discovery

After completing the preconfiguration process and creating a seed router, you can discover the TE network for a particular TE provider. This populates the repository with the network topology. Also, you might need to set up the management interfaces. The necessary steps are described in this section.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning the preconfiguration steps takes place.

Figure 8-6 Prime Provisioning Process Diagram - Preconfiguration

The purpose of the TE discovery process is to populate the repository with the TE topology, TE tunnels, explicit paths, and static routes to tunnels present in the live network.

The TE discovery process uses a seed device to discover the MPLS TE network topology using either Telnet or SSH. All the Traffic Engineering routers in the network should be accessible via their TE ID.

TE Discovery is a schedulable task that can be run once or on a periodic basis. Any inconsistencies between the repository and the network are reported in the Discovery log. The service state information is updated incrementally by logging tunnel in-use Label Switched Paths (LSPs) and updating the service request (SR) state.

This section includes the following:

- [TE Discovery Prerequisites and Limitations, page 8-12](#)
 - [Accessing TE Routers for TE Discovery, page 8-12](#)
 - [Memory Shortage on Large Networks, page 8-12](#)
 - [IOS XR and Enable Passwords, page 8-13](#)
- [Creating a TE Discovery Task, page 8-13](#)
 - [TE Incremental Discovery, page 8-13](#)
 - [TE Full Discovery, page 8-14](#)
- [Managing Per Area Discovery, page 8-15](#)
 - [Performing a Per Area TE Discovery, page 8-15](#)
 - [Running a Per Area TE Discovery Through an ABR, page 8-16](#)
- [Verifying a TE Discovery Task, page 8-16](#)
 - [Task Logs, page 8-16](#)
 - [TE Topology, page 8-19](#)
 - [View Network Element Types, page 8-19](#)
- [Setting Up Management Interfaces, page 8-19](#)
 - [MPLS-TE Management Process, page 8-19](#)
 - [Configuring Ethernet Links, page 8-19](#)

TE Discovery Prerequisites and Limitations

The following prerequisites apply mainly to TE discovery.

For an overview of the general Prime Provisioning prerequisites and limitations, see [Prerequisites and Limitations, page 8-3](#).

Accessing TE Routers for TE Discovery

To successfully run a TE discovery task, the seed router must be directly accessible from the management station.

All TE routers must be accessible from the Prime Provisioning machine via their TE router ID. This is often the loopback IP address, but not always.

For Telnet/SSH, there must be direct Telnet/SSH access from the Cisco Prime Provisioning Traffic Engineering Management (TEM) management station to each device.

See [Preconfiguration Process Overview, page 8-4](#) for instructions on how to select Telnet or SSH when setting up a seed router.



Note

After performing a TE discovery, it is recommended that you do not manually reconfigure RSVP graceful restart on the device. This affects the synchronization with the database and can cause deployment failure, in which case a new TE discovery needs to be performed.

Memory Shortage on Large Networks

When running TE Discovery on a large network (250+ devices or 5000+ tunnels, for example) or an `OutOfMemoryException` is encountered, it is recommended that the memory setting be changed.

To do this, use the following steps:

-
- Step 1** Choose **Administration > Hosts**.
 - Step 2** Select a host and click the **Config** button.
 - Step 3** Select **watchdog > server > worker > java > flags**.
 - Step 4** Change the first part of the property string, for example to **-Xmx1024m** instead of the default value **-Xmx512m**.

This increases the heap size of the **TE Discovery** task, which will clear up the `OutOfMemoryException` problem.
 - Step 5** Revert the **watchdog.server.worker.java.flags** property back to its original value to reduce the resource usage when no longer needed.
-



Note

Alternatively, the same memory increase can be achieved by editing the **watchdog.server.worker.java.flags** property in the **vpnsnc.properties** file.

IOS XR and Enable Passwords

If an IOS XR device is to be used as a seed device, the enable password should be set in its device record even though IOS XR does not require an enable password, for itself. That way IOS devices in the network, which do require an enable password, can be fully discovered.

When creating an IOS XR device through the **Devices** tab (**Inventory > Devices**) to act as a seed device for an initial discovery, it is not necessary to specify the enable password - TEM will be able to log in and get all the data it needs.

However, if there are other IOS devices in the same network, TEM will not be able to enter enable mode for those devices. As a result, these are not fully discovered in the sense that the inability to enter enable mode stops TEM from gathering all the relevant data. These other IOS routers will show up as **'unknown'** devices in the **Devices** window.

Limitations

Simultaneous TE Discovery in the same TE Provider is not supported. Only one user can run a TE Discovery per TE Provider at a time.

Creating a TE Discovery Task

In the Task Manager, you can run two types of TE Discovery tasks:

- [TE Incremental Discovery, page 8-13](#)
- [TE Full Discovery, page 8-14](#)

TE Incremental Discovery

This rediscovery process can take a long time to complete for a larger OSPF area.

In TE Incremental Discovery, the discovery tasks are run in increments whenever changes occur in the network, such as when a new device or link is added, causing a much smaller memory overhead than a TE Full Discovery.

To create a TE Discovery task on the TE network, use the following steps:

Step 1 Choose **Operate > Task Manager**.

The Task Manager window appears.

Step 2 Choose **Create > TE Incremental Discovery**.

The Task Creation wizard appears.

Step 3 Optionally, alter the **Name** and/or **Description** fields and click **Next**.

The TE Provider window appears.

Step 4 Select a TE provider and click **Next**.

The Device/Link Discovery Information window appears.

You can perform either of the following:

- Device discovery—A new device added to the network can be discovered using Device Discovery. For device discovery, non-Cisco devices, if any, are excluded from the list.

A device can be selected by clicking the **Select** button which shows the list of devices added in Inventory.

The prerequisite here is that the device which needs to be discovered needs to be added with its management IP address. The credentials of the device need not be the same as the credentials of other devices already populated in the repository. The device is successfully discovered only if it falls under the same OSPF area that is mentioned for the TE provider.

- **Link discovery**—A new link added to the network can be discovered using Link Discovery. Any explicit paths, primary, and backup tunnels traversing through that link will also be discovered.

End Device A and End Device B can be selected from the list of devices which have already been (TE Nodes). You must specify Interface A and Interface B.

Step 5 Select the seed device for discovering the network and click **Next**.

The Task Schedules window appears.

Step 6 Create a task schedule in one of two ways:

- Click **Now** to schedule the task to run immediately, in which case the schedule information is automatically filled into the Task Schedules list.
- Click **Create** to create a scheduler for this task, in which case the Task Schedule window appears.

Step 7 In the Task Schedule window, make your selections to define when and how often the task should be run.



Note

The default setting is to schedule a single **TE Discovery** task to take place immediately ("**Now**").

Step 8 Click **OK**.

The scheduled task should now appear in the Task Schedules table.

Step 9 Click **Next**.

A summary of the scheduled task appears.

Step 10 Click **Finish**.

This will add the task to the list of created tasks in the Tasks window.

TE Full Discovery

In a TE Full Discovery, the discovery task runs without stopping until all devices have been discovered. To create a TE Discovery task on the TE network, use the following steps:

Step 1 Choose **Operate > Task Manager**.

The Task Manager window appears.

Step 2 Create a new task by selecting **Create > TE Full Discovery**.


The Create Task window appears.

Step 3 Optionally, alter the **Name** and/or **Description** fields and click **Next**.

The Select TE Provider window appears.

Step 4 Select a TE provider and click **Next**.

The Select Seed Device window appears. Non-Cisco devices, if any, are excluded from the list.

- Step 5** Select the seed device for discovering the network and click **Next**.
The Task Schedules window appears.
- Step 6** Create a task schedule in one of two ways:
- Click **Now** to schedule the task to run immediately, in which case the schedule information is automatically filled into the Task Schedules list.
 - Click **Create** to create a scheduler for this task, in which case the Task Schedule window appears.
- Step 7** In the Task Schedule window, make your selections to define when and how often the task should be run.
-  **Note** The default setting is to schedule a single **TE Discovery** task to take place immediately (“**Now**”).
- Step 8** Click **OK**.
The scheduled task should now appear in the Task Schedules table.
- Step 9** Click **Next**.
A summary of the scheduled task appears.
- Step 10** Click **Finish**.
This will add the task to the list of created tasks in the Tasks window.

Managing Per Area Discovery

Before running a per area TE discovery, it is helpful to understand how multiple OSPF areas are managed by Prime Provisioning.

For background information on this topic, see the section Multiple OSPF Areas in [Traffic Engineering Management Concepts, page 8-111](#).

This section describes the following:

- [Performing a Per Area TE Discovery, page 8-15](#)
- [Running a Per Area TE Discovery Through an ABR, page 8-16](#).

Performing a Per Area TE Discovery

When a TE Discovery is run against an area with a selected TE provider, all tunnels and explicit paths associated with that area will be imported into the Prime Provisioning database.

To initiate a per area TE discovery, use the following steps:

- Step 1** Create an Provider.
- Step 2** Create an Region.
- Step 3** Create a TE Provider.
- Step 4** Create a seed device from the Devices window.
- Step 5** Choose **Operate > Task Manager > Create > TE Full Discovery**.
Specify a name for the TE Discovery task or accept the default and click **Next**.

- Step 6** Select a TE Provider and click **Next**.
- Step 7** Select a seed device and click **Next**.
- Step 8** Select a schedule for the TE Discovery and click **Next**.
- Step 9** Review the summary of the discovery task.
- If it is acceptable, click **Finish** to start the TE Discovery process.

Running a Per Area TE Discovery Through an ABR

If no area identifier is specified in the TE provider configuration and the seed device is an ABR, TE Discovery will abort with the warning message shown in [Figure 8-7](#) informing you to either specify an area identifier for the TE provider or use a non-ABR device as the seed.

Figure 8-7 *TE Discovery Through an ABR with no TE Area Identifier Specified*

Task Log			
		Log Level: Warning	Component: * Filter
Date	Level	Component	Message
2011-03-08 07:49:42	WARNING	repository.rbac	Thread RBAC enabled flag is set to false.
2011-03-08 07:49:55	SEVERE	DiscoveryTask	Seed device 192.168.1.139 has TE enabled in multiple IGP areas. This configuration is unsupported with the specified TE Provider, aborting discovery. Retry discovery from a seed device with TE enabled in one IGP area or specify the area you wish to be discovered by editing the TE Provider.
2011-03-08 07:49:55	WARNING	DiscoveryTask	Fatal Error Encountered, aborting Discovery...
2011-03-08 07:49:55	SEVERE	DiscoveryTask	Discovery FAILURE.
2011-03-08 07:49:55	WARNING	repository.rbac	Thread RBAC enabled flag is set to true.
			Return to Logs

Verifying a TE Discovery Task

The result of running the **TE Discovery** task can be assessed in four ways:

- **Task Logs**—View a summary log of any changes that have occurred in the network.
- **TE Topology**—Display the latest TE Topology from the repository.
- **View Network Element Types**—In the Traffic Engineering Management GUI, go to **TE Nodes**, **TE Links**, **TE Primary Tunnels**, and so on to verify the state of specific network element types.
- Viewing the state of discovered devices—Go to the Service Requests window to examine whether the state of the discovered devices is as expected.

Task Logs

The TE Discovery log captures the state of the network and compares it with the most recent snapshot of the repository.

To view the task log for a **TE Discovery** task, use the following steps:

- Step 1** Choose **Operate > Task Logs**.
- The Task Logs window appears.

The status of the task is shown in the **Status** column. This updates automatically and indicates when the TE Discovery process is complete.

If the task is not completed and **Auto Refresh** is selected, the table continues to update periodically until it is completed.

Step 2 To view the log for a particular task, go to **Operate > Task Manager**, select the desired task, and then click the **View Log** button.

A copy of a TE Discovery log is shown in the following screenshots, starting with [Figure 8-8](#). This first example shows the TE-enabled devices and links that TE Discovery has found in the topology. Once each device is identified, a set of debug, informational, warning and error logs are built up for each device to facilitate identification of errors.



Note To find the summary of changes in the network depicted in the following screenshots, scroll to the bottom of the log.

Figure 8-8 TE Discovery Task Log - Example 1

Task Log

		Log Level: All		Component: *	Filter
Date	Level	Component	Message		
2011-11-07 16:29:00	WARNING	repository.rbac	Thread RBAC enabled flag is set to false.		
2011-11-07 16:29:00	INFO	DiscoveryTask	Thread-specific rbac checking is turned off		
2011-11-07 16:29:00	INFO	DiscoveryTask	Provider: teprovider		
2011-11-07 16:29:00	INFO	DiscoveryTask	Seed Router: SOLKTXES8AW		
2011-11-07 16:29:00	INFO	DiscoveryTask	INFO: MplsTeDiscoveryHandler: customer set to: teprovider-default-customer		
2011-11-07 16:29:00	INFO	DiscoveryTask	INFO: MplsTeDiscoveryHandler: region set to: region		
2011-11-07 16:29:00	CONFIG	DiscoveryTask	DEBUG: fetching topology from seed device.		
2011-11-07 16:29:00	CONFIG	DiscoveryTask	DEBUG: successfully retrieved topology from seed device.		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.103		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.236		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.7		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.104		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.253		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.6		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.101		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.252		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.9		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.102		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.233		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.8		
2011-11-07 16:29:00	INFO	DiscoveryTask	MplsTeDiscoveryHandler: INFO: Found device in network, MPLS TE ID: 69.82.254.232		

[Figure 8-9](#) and [Figure 8-10](#) show a sample device debug and information section.

Figure 8-9 *TE Discovery Task Log - Example 2*

```

2011-11-07 16:47:30 INFO DiscoveryTask <----->
Information summary for Te Router, Te Id : , Host name: WJRDUT307AW
-
- NEW: Te Router created, Mpls Te Id: 69.82.254.103
-
- Device Interfaces:
-
- EXISTING: Interface found with no changes, Name: MgmtEth0/RP0/CPU0/0, IP Address: 10.141.218.17
- EXISTING: Interface found with no changes, Name: TenGigE0/4/3/0, IP Address: 69.82.120.81
- EXISTING: Interface found with no changes, Name: Loopback10, IP Address: 10.214.254.103
- EXISTING: Interface found with no changes, Name: MgmtEth0/RP1/CPU0/0, IP Address: 10.141.218.18
- EXISTING: Interface found with no changes, Name: TenGigE0/4/1/0.1100, IP Address: 69.82.122.140
- EXISTING: Interface found with no changes, Name: TenGigE0/3/3/0, IP Address: 69.82.120.79
- EXISTING: Interface found with no changes, Name: Loopback0, IP Address: 69.82.254.103
- EXISTING: Interface found with no changes, Name: TenGigE0/4/4/0.1100, IP Address: 69.82.122.142
- EXISTING: Interface found with no changes, Name: TenGigE0/13/0/0, IP Address: 69.82.122.134
- EXISTING: Interface found with no changes, Name: TenGigE0/10/0/0, IP Address: 69.82.122.132
- EXISTING: Interface found with no changes, Name: TenGigE0/4/4/0.1250, IP Address: 69.82.77.128
- EXISTING: Interface found with no changes, Name: TenGigE0/3/0/0, IP Address: 69.82.77.48
- EXISTING: Interface found with no changes, Name: TenGigE0/4/1/0.1250, IP Address: 69.82.77.132
- EXISTING: Interface found with no changes, Name: TenGigE0/4/2/0, IP Address: 69.82.77.54
- EXISTING: Interface found with no changes, Name: TenGigE0/3/2/0, IP Address: 69.82.77.52
- EXISTING: Interface found with no changes, Name: TenGigE0/4/0/0, IP Address: 69.82.77.50
-
- Te Links:
-
2011-11-07 16:47:30 CONFIG DiscoveryTask <----->
Debug summary for Te Router, Te Id: 69.82.254.236, Host name: TWBGOHAA81W
DEBUG: Calling device for show version output: 69.82.254.236

```

269255

Figure 8-10 *TE Discovery Task Log - Example 3*

```

2011-11-07 16:47:30 CONFIG DiscoveryTask <----->
Debug summary for Te Router, Te Id: 69.82.254.103, Host name: WJRDUT307AW
-
DEBUG: Calling device for show version output: 69.82.254.103
DEBUG: MplsTeShowVersionCallback: XDE show version invocation completed normally for device:
69.82.254.103
DEBUG: MplsTeShowVersionCallback: Device: 69.82.254.103, has an OS with version: 4.0.1[Default]
DEBUG: MplsTeShowVersionCallback: Device: 69.82.254.103, is running Cisco IOS XR.
DEBUG: Calling device for show running-config output: 69.82.254.103
DEBUG: Calling device for show primary tunnels output: 69.82.254.103
DEBUG: Calling device for show backup tunnels output: 69.82.254.103
DEBUG: MplsTeShowRunningCallback: XDE show running config invocation , MPLS TE ID:
69.82.254.103, completed normally.
DEBUG: MplsTeShowRunningCallback: Device has the following flags: rsvp graceful restart: false, te
enabled: true, conformant: true, supports FRR true, snmp traps enabled: true
DEBUG: Calling device for show auto-bw output: 69.82.254.103
DEBUG: MplsTeShowTunnelsCallback: show tunnels command completed successfully on device:
69.82.254.103, found tunnels: 1000 1001 1003 1004 1005 1006 1008 1009 1010 1013 1014 1017 1020
1023 1024 1025 1028 1029 10100 10101 10200 10201 10300 10301 10400 10401 10500 10501
10600 10601 10700 10701 10800 10801 10900 10901 11000 11001 11100 11101 11200 11201 11400
11401 11500 11501 11600 11601 11700 11701 11800 11801 11900 11901 12100 12101 12300 12301
12500 12501 12700 12701 14100 14101 14200 14201 15800 15801 16100 16101 16200 16201 16300
16301 16400 16401 18100 18101 18200 18201
DEBUG: Calling device for show supports subpool output: 69.82.254.103
DEBUG: MplsTeShowTunnelsBackupCallback: show backup tunnels command completed successfully
on device: 69.82.254.103, found backup tunnels: 1000 1001 1003 1004 1005 1006 1010 1013 1014
1017 1020 1023 1024 1025 1028 1029
DEBUG: MplsTeShowAutoBwCallback: XDE show auto bw invocation for device, MPLS TE ID:
69.82.254.103, completed normally.
DEBUG: MplsTeShowAutoBwCallback: Device: 69.82.254.103, supports auto bandwidth.
DEBUG: MplsTeShowSubpoolCallback: show supports subpool command completed successfully on
device: 69.82.254.103
DEBUG: MplsTeShowSubpoolCallback: this device supports subpool.
DEBUG: MplsTeShowRunningCallback: Device: WJRDUT307AW, has TE enabled interfaces:
TenGigE0/4/3/0, TenGigE0/4/1/0.1100, TenGigE0/3/3/0, TenGigE0/4/0/0.1100, TenGigE0/13/0/0,
TenGigE0/10/0/0
Device: WJRDUT307AW, has non TE enabled interfaces: MgmtEth0/RP0/CPU0/0, Loopback10,
MgmtEth0/RP1/CPU0/0, Loopback0, TenGigE0/4/4/0.1250, TenGigE0/3/0/0, TenGigE0/4/1/0.1250,
TenGigE0/4/2/0, TenGigE0/3/2/0, TenGigE0/4/0/0
DEBUG: MplsTeShowRunningCallback: Device: WJRDUT307AW, has explicit paths: WJRDUT307AW-
AURSCOTY7AW-1 WJRDUT307AW-AURSCOTY7AW-3 WJRDUT307AW-CLSPCOYK8BW-1 WJRDUT307AW-
HCHLIMT7AW-2 WJRDUT307AW-HLBOOR387AW-1 WJRDUT307AW-HLBOOR387AW-2
WJRDUT307AW-OMALNEXU7AW-4 WJRDUT307AW-RCKLCAIG7AW-1 WJRDUT307AW-
RCKLCAIG7AW-2 WJRDUT307AW-RCKLCAIG7AW-3 WJRDUT307AW-RCKLCAIG8AW-3
WJRDUT307AW-RCKLCAIG8AW-4 WJRDUT307AW-RCKLCAIG8BW-1 WJRDUT307AW-
RCKLCAIG8BW-2 WJRDUT307AW-RCKLCAIG8BW-3 WJRDUT307AW-RDMWA227AW-1
WJRDUT307AW-RDMWA227AW-3 WJRDUT307AW-RDMWA227AW-4 WJRDUT307AW-
SCRMCAGN81W-1 WJRDUT307AW-SOLKTXES8AW-2 WJRDUT307AW-SOLKTXES8BW-1
DEBUG: MplsTeShowRunningCallback: Device: WJRDUT307AW, has tunnels: 1003 1004 1001 10500

```

269256

Step 3 Click **Return to Logs** to quit the current log with the option to open another log.

TE Topology

The TE Topology tool provides a visual snapshot of the current state of the network. It cannot be used to determine changes that have taken place in the network.

The steps required to generate a topology graph of the network are described in [TE Topology, page 8-80](#).

View Network Element Types

Another way to check the state of the network after running TE discovery is to go to the Traffic Engineering menu options and select the type of elements you want to verify.

For example, to check the status of the nodes after running TE discovery, choose **Traffic Engineering > Nodes**. Look at the updated list of TE nodes to assess which nodes are in the network.

Do the same for TE Links, TE Primary Tunnels, TE Backup Tunnels, and so on.

Setting Up Management Interfaces

Before commencing tunnel management operations, you need to set up management interfaces. However, this step is only necessary if the network devices are not accessible by the hostname from the management station.

For a detailed description of how to set up management interfaces on specific devices, see [Devices, page 2-1](#).

MPLS-TE Management Process

The MPLS-TE management process involves the following steps:

1. Enable MPLS-TE on the network devices and make sure that the IP addresses used as the devices TE IDs are accessible from the management station (this step is not supported by TEM).
2. Prepare the repository for discovering MPLS-TE network.
3. Set up management interfaces for the discovered devices or update the server host file with resolution for all discovered devices. Again, this is not needed if the hostnames are already accessible from the management station.
4. Discover the MPLS-TE network.

You will then be in a position to run the other MPLS-TE functions available in TEM.



Note

When the repository is empty, or when the management IP addresses are not configured for current devices in the TE network, make sure that the router MPLS TE ID can be reached from the management station. In other words, the TE discovery process does not support seed passthrough.

Configuring Ethernet Links

Only point-to-point links are supported in TEM. POS links are point-to-point by default but otherwise Ethernet links need to be configured as point-to-point.

For IOS, enter the following command:

```
(config-if)# ip ospf network point-to-point
```

For IOS XR, enter the following command:

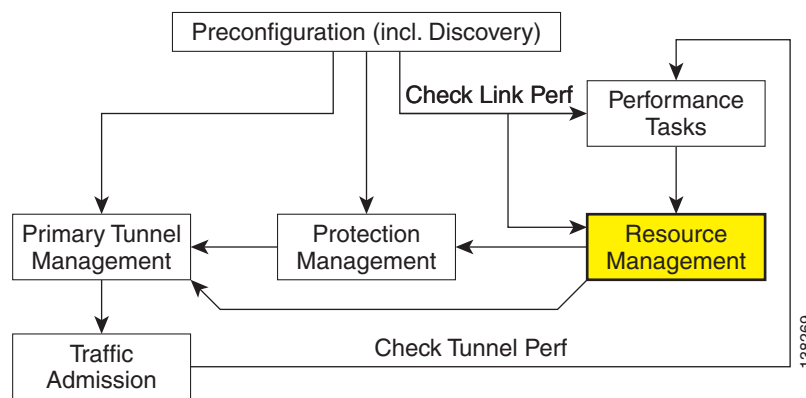
```
# router ospf <id> area <area identifier> interface <name> network point-to-point
```

TE Resource Management

TE resource management is defined as the tuning of certain properties on the TE interfaces to optimize the tunnel placement.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning resource management occurs.

Figure 8-11 Prime Provisioning Process Diagram - Resource Management

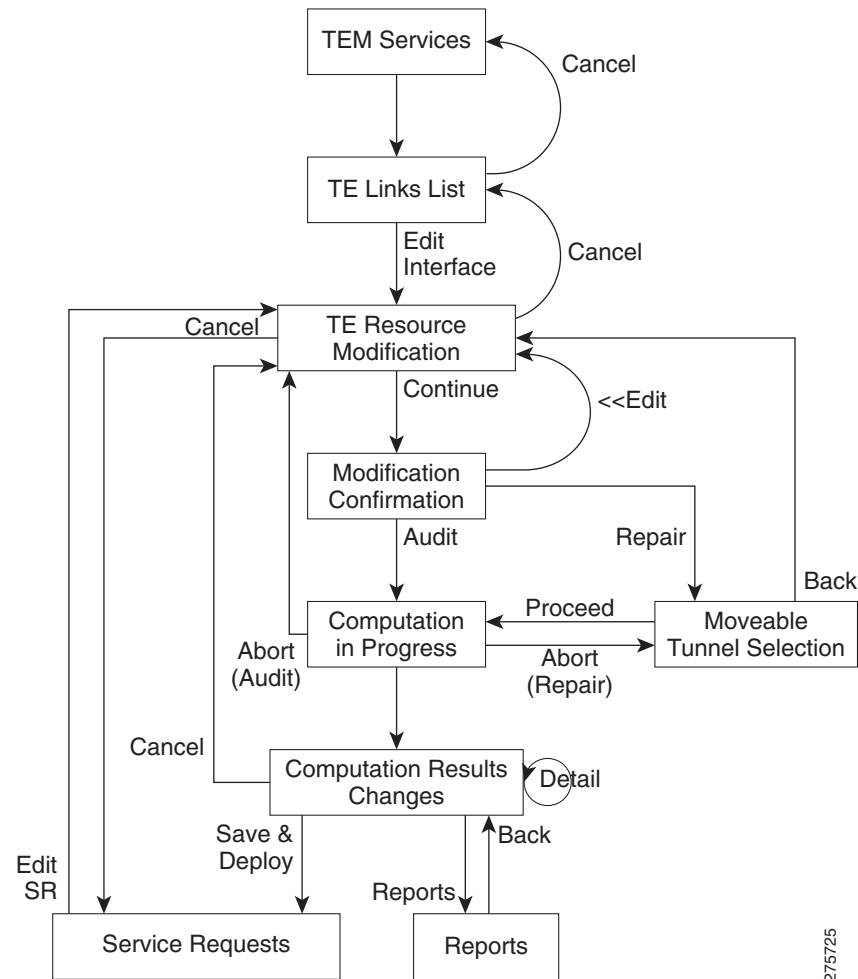


When a tunnel placement is attempted and there is insufficient bandwidth, sometimes the resources on the TE links can be changed and the tunnel placement retried.

Network resources in this context are understood to be routers in the TE network, the interfaces that connect them, and the RSVP bandwidths and other properties configured on the links. Because Prime Provisioning relies on the discovery process to add the network elements to the repository, the resources must be discovered before resource management can be performed.

TE resource management is a manual process that should be performed on an as needed basis. If the original configuration is already optimal, there is no need to do any resource management tasks. If subsequent discovery unveils any discrepancy, or if you experience difficulty achieving desired results in protection planning or placing primary tunnels, adjustments on the resources might be warranted.

An overview of the resource management process is provided in [Figure 8-12](#).

Figure 8-12 Resource Management Processes

This section includes the following:

- [Modifying Network Resources, page 8-21](#)
- [Changing Link Status, page 8-23](#)
- [Deleting TE Links, page 8-24](#)
- [Deleting TE Tunnels, page 8-25](#)
- [Deleting TE Nodes, page 8-26.](#)

Modifying Network Resources

The resource management tasks are mainly carried out from the TE Links List window.



Note

Certain attributes, such as Description, that do not impact the computation carried out by these tools and updates to these are, therefore, not displayed in the computation results window.

To modify a TE link, use the following steps:

Step 1 Choose **Traffic Engineering > Links**.

The TE Links List window appears.

The links list shows the current active links in the TE network. Use the arrows to page forward as needed.

Step 2 Select the desired link in the links list.

Note **Admin Status**—Indicates whether the link is **UP** or **DOWN**. This is local to Prime Provisioning. It is not the network interface status.

Step 3 Click **Edit > Interface A** or **Edit > Interface B** to edit one of interfaces on the link.

Note If a non-Cisco interface is selected for editing, changes made in the Edit window will be saved in the Prime Provisioning repository but they will not be deployed.

The TE Resource Modification window appears. It includes the following fields:

- **Max Global (BC0) Reservable**—Maximum amount of bandwidth in kbps that can be reserved by TE Tunnels.
- **Max Sub Pool (BC1) Bandwidth**—Maximum amount of bandwidth in kbps that can be reserved by sub pool TE Tunnels. The range is from 1 to the value of **Max Global Reservable**.
- **Attribute Bits**—Links attributes to be compared to a tunnel's affinity bits during selection of a path. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits) where the value of an attribute is 0 or 1.
- **TE Metric**—Metric used to override the Interior Gateway Protocol (IGP) administrative weight (cost) of the link.
- **Propagation Delay**—The time it takes for traffic to travel along a link from the head interface to the tail interface.
- **Max Delay Increase**—Used in computations of FRR backup-tunnels to constrain the propagation delay of a backup-tunnel for the link. A max delay increase for a link might need to be set to loosen the delay constraint when generating backup tunnels, as it is difficult to find backup tunnel paths where there is no increase in the delay compared with the flow being protected.
- **Link Speed Factor**—Multiplication factor corresponding to the amount (percentage) of link speed available for primary and backup traffic. This is typically set to 1.

Step 4 Make the desired modifications and click **Continue** to proceed to the confirmation page to verify the changes or click **Cancel** to quit without saving.**Step 5** Click **Edit** to return to the editable window or proceed in one of the following ways:

- **Proceed with Changes** —Perform Tunnel Audit or Tunnel Repair.

For a detailed explanation of Tunnel Audit and Tunnel Repair, see [Advanced Primary Tunnel Management, page 8-44](#)

If a non-Cisco device is edited, **Proceed with Changes** will be disabled. Instead, **Save & Deploy** is enabled and the changes can be saved (not deployed).

- **Save & Deploy**—If the changes made do not affect tunnel placement, click **Save & Deploy** to proceed. In this case, there is no need for performing Tunnel Audit or Tunnel Repair.

**Note**

When you click **Save & Deploy**, a background process is started. To avoid a potential conflict with another deployment, wait until the service request (SR) has completed the Requested and Pending states before deploying another SR with Save & Deploy. To see the state of deployment, go to **Operate > Service Request Manager** or open **Operate > Task Manager**.

**Note**

In Prime Provisioning, service requests (SRs) are generally deployed from each TE service, not from the **Operate > Service Request Manager** page with the exception of the TE Traffic Admission SR.

After deployment, the SR status can be viewed from the SR window at **Operate > Service Request Manager**.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

If the SR does not go to the **Deployed** state, go to the Task Log to see the deployment log (**Operate > Task Manager > Logs**). Task logs are further described in [Task Logs, page 8-16](#).

Changing Link Status

From the TE Links List window, you can also find out what effect it will have if a link is taken offline. This approach can be used to move tunnels off a link before actually shutting down the interface.

**Note**

Link status in Prime Provisioning is of local significance. Changing link status as described in this section is not provisioned down to the network.

To change the link status, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Links**.
- The TE Links List window appears.
- Step 2** Select one or more links and click the **Change Status** button.
- Step 3** Select **Enable** or **Disable** to enable or disable the selected link.
- As an example, selecting **Disable** will change the link status to **DOWN**. Similarly, use **Enable** to change the status back to **UP**.
- Step 4** Click **Proceed with Changes** to assess any impact on tunnel placement using Tunnel Audit or Tunnel Repair and deploy the changes.
- For a detailed explanation of Tunnel Audit and Tunnel Repair, see [Advanced Primary Tunnel Management, page 8-44](#).
-

Deleting TE Links

The TE Link List window includes a delete function (the **Delete** button), which allows you to delete a TE link and the TE interfaces at each end of the link from the repository. It does not make any change to the physical link in the network.

Link deletion can be selected based on a specific TE provider. When deleting different links belonging to different providers, first choose the appropriate provider and then mark the links to be deleted.

Also, simultaneous deletion of multiple links of the same provider is supported.

Restrictions

The Prime Provisioning GUI prevents you from deleting a link if any TE object is still using that link.

It checks the following objects:

- strict explicit paths
- protected interfaces of backup tunnels
- SRLGs
- protected elements
- TE resource SRs.

If there are any primary or backup tunnels traversing the path options, an error report will be displayed. Otherwise, a message will be displayed seeking confirmation that the above set of associated objects should be deleted.

Use Case

In this example, we will look at the procedure required when attempting to delete a link that could be traversed by primary or backup tunnels.

Use the following steps:

-
- Step 1** Choose **Traffic Engineering > Links**.
 - Step 2** Select a link by checking the corresponding check box.
 - Step 3** Click the **Delete** button.
 - Step 4** Two things can happen:
 - A tunnel with path option traverses the link: The link deletion will fail and you will be prompted to reroute or delete those tunnels before trying link deletion again. This will take you to the TE Links List page.
 - No tunnels with path option traverses the link: A list of TE associated objects will be displayed for that link and you will be prompted to confirm whether you agree to the automatic deletion of TE Link associated objects or have second thoughts and would like to cancel the link deletion transaction.
 - Step 5** After any necessary tunnels have been rerouted/deleted and link deletion is attempted, a list of objects that are still associated will be displayed.

- Step 6** If you want to delete associated TE objects listed after rerouting/deleting primary tunnels, you will get directed to a new window showing the progress of the transaction only when there are tunnels offering backup link protection/protecting multiple interfaces. If there are no tunnels offering backup link protection/protecting multiple interfaces, you are directed to the TE Links window on successful/failure transaction from the associated TE objects list page.
- See the note below on associated TE objects.
- Step 7** After all the associated objects have been deleted, you will be directed to the TE Links List window.
-

Note on Associated TE Objects

Associated TE objects can be any of the following:

- strict explicit paths and loose explicit paths (with strict hop type) traversing the link;
- backup tunnels offering link protection;



Note The link will be removed from any SRLGs (if SRLG has more than one link) or both the link and the SRLG will be removed if the link marked for deletion is the only one in the SRLG.

- resource SRs;
- protected elements.

The associated TE objects in the above list vary depending on the way the link is configured in TEM.

For example, if associated TE objects have backup tunnels offering link protection, you will be directed to the Link Deletion Progress window where protected interfaces will be updated accordingly for the available TE links and backup tunnel SRs will get re-deployed. Otherwise, if no backup tunnels offering link protection qualify as associated TE objects, the remaining TE objects will automatically be removed from the window showing the associated TE objects.

Deleting TE Tunnels

TE Tunnels can be deleted in the TE Links List window or in the individual primary or backup tunnel SR windows (see [Delete Primary Tunnel, page 8-38](#) or [Delete Backup Tunnel, page 8-43](#)).

In the TE Links window, the reason for wanting to delete a tunnel will often be a need to delete a link that is traversed by one or more tunnels.

To delete a tunnel in the TE Links List window, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Links**.
- Step 2** Select the link for the tunnel that you wish to delete and click the **Show Tunnels** button.
- This brings up a tunnel filter where you can select the category of tunnel you wish to display (**All**, **Managed**, **Unmanaged**, **Backup**).
- Step 3** Select one of these tunnel categories.
- This brings up a list of all tunnels in the selected filter category, which traverses the link.
- Step 4** Select one or more tunnels that you wish to delete and click the **Delete** button.

This will delete the tunnels selected by starting a new provisioning operation.

Deleting TE Nodes

You can also delete a TE node. This works in a very similar way to deleting a link but is done from the PE devices screen. By deleting the corresponding PE device, you effectively delete the TE node.

Similar restrictions apply as in the case of TE links. The delete operation can only be succeed if no TE objects are using the node.

Restrictions

The Prime Provisioning GUI prevents you from deleting a node if any TE object is still using that node.

As with TE links, it checks the following objects:

- strict explicit paths
- protected interfaces of backup tunnels
- SRLGs
- protected elements
- TE resource SRs.

In addition, the node deletion checks that no managed, unmanaged, or backup tunnel starts or ends at the node in question.

If any of these objects is using the node, an attempt to delete the node will result in an error message and the node and its interfaces remain unchanged.

Use Case

An example of this feature is when a TE router is to be decommissioned from the network and replaced by one or more new TE routers as part of a major topology change.

The steps needed to enable you to delete this node might include the following:

1. Reroute all managed tunnels away from this node using Tunnel Repair.
2. Reroute all unmanaged and backup tunnels using the node as part of their path away from it.
3. Delete any backup tunnels that protect either of the interfaces that make up the node.
4. Delete any explicit paths that use the node.
5. Delete the node from the repository from the TE Links List window.
6. Outside Prime Provisioning, during a suitable outage window, physically decommission the node, and set up its replacement(s).
7. Run a new TE discovery task, which result in the newly added nodes being added to the repository.
8. Depending on the FRR requirements of the network, protect the new node(s) using Compute Backup. (See [Compute Backup](#), page 8-63.)
9. Run network grooming (see [Grooming](#), page 8-57) to optimise the managed tunnels, so that they will make use of the new node(s).

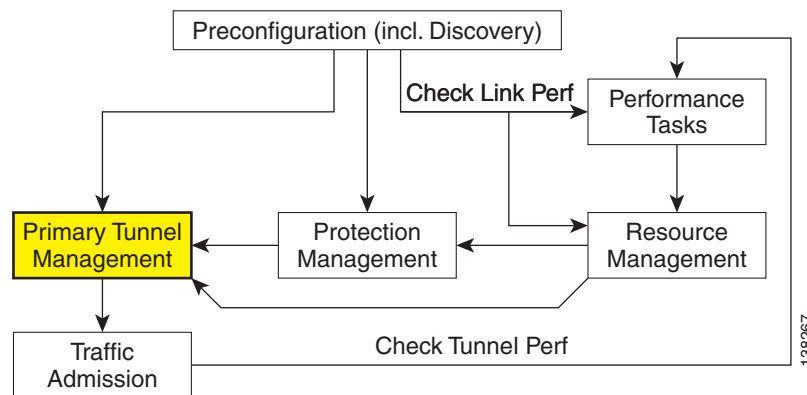
If this check succeeds, the TE node and all TE links and TE interfaces starting at that node are removed from the repository.

Basic Tunnel Management

This section describes the processes involved in creating primary and backup tunnels with Prime Provisioning. To create a tunnel, certain steps must first be performed as described in previous sections.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning primary tunnel management occurs.

Figure 8-13 Prime Provisioning Process Diagram - Primary Tunnel Management



Primary tunnels are characterized by carrying traffic during normal operation. They have a prioritized list of possible paths, by which traffic can be routed. At any one time, the highest priority path available will be used to route traffic. If this fails, traffic will normally be rerouted via the next available path until a higher priority path becomes available again.

Prior to setting up the tunnel, a TE policy governing the traffic must be defined. An explicit path is created to establish the route and, in the case of a primary tunnel, it is created as either a managed or an unmanaged tunnel.

The purpose of a backup tunnel is to carry Fast Re-Route (FRR) protected traffic around a failed element until the routing in the network has reconverged. It is intended to protect traffic travelling along primary tunnels. There can be many backup tunnels protecting the same traffic through the use of load balancing.

If the network fails to reconverge, the backup tunnel will remain in place.

The difference between managed and unmanaged tunnels is described in the section on Managed/Unmanaged Primary Tunnels in [Traffic Engineering Management Concepts, page 8-111](#).

The concept of bandwidth pools from which tunnels reserve bandwidth is important to understand. This is described in the section on Bandwidth Pools in the [Traffic Engineering Management Concepts, page 8-111](#).

This section includes the following:

- [Create TE Policy, page 8-28](#)
- [Create Explicit Path, page 8-29](#)
 - [Delete Explicit Path, page 8-31](#)

- [Primary Tunnel Operations, page 8-31](#)
 - [Create Primary Tunnel, page 8-32](#)
 - [Edit Primary Tunnel, page 8-37](#)
 - [Delete Primary Tunnel, page 8-38](#)
- [Backup Tunnel Operations, page 8-39](#)
 - [Create Backup Tunnel, page 8-39](#)
 - [Edit Backup Tunnel, page 8-42](#)
 - [Delete Backup Tunnel, page 8-43](#)

Create TE Policy

To create a primary tunnel, each primary tunnel must be associated with a policy. A policy can be used by multiple tunnels.

For backup tunnels, this step is not necessary. In this case, proceed to [Create Explicit Path, page 8-29](#).

For other TE policy management operations, see [TE Policies, page 8-71](#).

The TE policy is a set of rules governing the TE network and defines the Class-of-Service (for example, gold, silver, bronze) for primary tunnel traffic.

Prime Provisioning has a notion of **Managed** and **Unmanaged** policies. **Managed** policies have setup/hold priorities of 0/0 and can have additional routing constraints such as protection level and max delay. Tunnels with **Unmanaged** policies are provisioned by the system, but the system only tracks the deployment, not the operation of the tunnel. **Unmanaged** policies cannot have a setup/hold priority of zero.

For more information about managed and unmanaged primary tunnels, see the section on Managed/Unmanaged Primary Tunnels in the [Traffic Engineering Management Concepts, page 8-111](#).

Policies are managed under **Policies in Service Design**. For a more detailed explanation of the **Policies** GUI, see [TE Policies, page 8-71](#).

To create a TE policy, use the following steps:

Step 1 Choose **Traffic Engineering > Policy Manager**.

The Policy Manager window appears.

Step 2 Click **Create** and select **TE Policy** to set up a new TE policy.

To edit an existing policy, select the policy that you want to modify and click **Edit**. The TE Policy Editor window appears.



Note A policy that is being used by a tunnel cannot be modified. However, the name and ownership of an in-use policy can be changed.

For an explanation of the various window elements, see [TE Policies, page 8-71](#).

Step 3 Fill in the required fields marked with an asterisk (*) and any optional fields.

If you intend to use the TE policy for managed tunnels, make sure to check the **Managed** check box.

When setting up a policy for a managed tunnel, the **Setup** and **Hold** priorities are automatically set to zero (highest priority). In the case of a policy for an unmanaged tunnel, you can specify the desired **Setup** and **Hold** priority settings.

Step 4 Click **Save**.

Create Explicit Path

Paths are defined between source and destination routers, possibly with one or more hops in between. Paths are used for primary and backup tunnels in the explicit path option(s).

If you intend to create an explicit path for managed tunnels, the path should not contain any non-TE enabled interfaces. Paths with non-TE enabled interfaces will be filtered out by the tunnel path chooser of the tunnel editor for managed tunnels and backup tunnels (not unmanaged tunnels).

To create or edit an explicit path, use the following steps:

Step 1 Choose **Traffic Engineering > Explicit Paths**.

The TE Explicit Path List window appears.

Step 2 To create an explicit path in the **TE Explicit Path List**, click **Create**.

The New TE Explicit Path window appears.

To edit an explicit path in the explicit path list, select the explicit path that you want to modify and click **Edit**. This opens the TE Explicit Path Editor window.



Note An explicit path that is being used by a tunnel cannot be modified. However, use Edit to view the path.

The New TE Explicit Path window includes the following GUI elements:

- **Path Name**—Name of explicit path.
- **Head Router**—Name of the head router.
- **Path Type**—Three types of explicit paths are supported:
 - **STRICT**—All strict hops are defined in the path.
 - **LOOSE**—Any loose hops (pure loose path or a combination of loose and strict hops) are defined in the path.
 - **EXCLUDE**—All exclude hops are defined in the path.
- **Links** (table)—Lists the links added for the current path and includes the following information:
 - **Device**—Hostname of the TE device that the path originates from.
 - **Outgoing Interface**—Interface name of the outgoing interface from the originating device.
 - **Outgoing IP**—IP address of the outgoing interface.
 - **Next Hop**—Hostname of the next hop device.
 - **Incoming Interface**—Incoming interface name on the next hop device.
 - **Incoming IP**—Incoming interface IP address on the next hop device.

- **Provision Preference**—Preference for provisioning the **next-address** subcommand of the **ip explicit-path** command. Choose between **Outgoing Interface** and **Incoming Interface**.
 - **Outgoing Interface**—Outgoing interface on the router.
 - **Incoming Interface**—Incoming interface on the router.



Note If a path is used by any tunnel, no modifications are possible. The **Outgoing Interface** and **Incoming Interface** links are not selectable and the Provision Preference line and the **Add Link**, **Delete Link**, and **Save** buttons disappear.

Step 3 Specify a pathname and select a head router.

Step 4 Select a path type:

- **Strict**: If **Strict** is chosen, use the current panel that lists the connected links one by one until destination is reached.
- **Loose**: If **Loose** is selected, a new hop is added by entering the IP address. If **Strict** is selected, you are allowed to select from TE Links list only.



Note For IOS XR, the **Loose** type is only available if the head device is running IOS XR 3.4 or later.



Note If **Loose** is chosen, a new panel that adds a loose hop definition one by one is listed. Because a combination of strict and loose hops is allowed for a loose explicit path definition, the flexibility of including strict hops is provided with a constraint of at least a loose hop presence in the path.

- **Exclude**—**Exclude** allows you to specify an exclude IP address. See [Step 6](#).

Step 5 If **Strict** was selected, click the **Add Link** button to add a blank line to the hop list table.

If **Loose** or **Exclude** was selected, an **Add Hop** button appears, which when clicked opens a pop-up window where you specify an IP address.

Step 6 Now an interface must be selected for the head router.

Depending on the path type selection, you will see one of the following windows:

A. Strict path type:

Click the **Add Link** button, then click **Add Interface**. The Select Next Hop window appears.

The next hop list contains all the possible next hops of the router, excluding the ones already included in the explicit paths (to avoid path loops).

The next hop list contains TE interfaces and at most one non-TE interface for each router (if the loopback interface is used as the MPLS TE ID of the device). For TE interfaces, the **Outgoing Interface** and **Outgoing IP** columns are populated by the application.



Note If a non-TE interface is selected, **Provision Preference** is set to **Incoming Interface**. The provision preference cannot be set manually.

Select an interface and click **Select**. The corresponding link information is added to the new explicit path in the **Links** table.

In the New TE Explicit Path window, both the incoming and outgoing interface fields are populated.

B. Loose path type:

Click the **Add Hop** button. The Loose Hop Definition window appears.

In this window, specify an IP address for the desired loose hop and click **OK**. The Loose Hop Definition window closes.

The New TE Explicit Path window now displays the added loose hop.

C. Exclude path type:

Click the **Add Hop** button. The Exclude Hop Definition window appears.

In this window, specify an IP address for the desired exclude hop and click **OK**. The Exclude Hop Definition window closes.

The New TE Explicit Path window now displays the added exclude hop.

Step 7 To add another link, click either **Add Link** or **Add Hop**.

Step 8 For Strict hops, a **Provision Preference** can optionally be selected by clicking either the **Outgoing Interface** or the **Incoming Interface** radio button.



Note If you try to select the **Provision Preference** before adding a link when non-TE interfaces are present, the **Add Link** process overrides the **Provision Preference** and sets it to incoming.

Step 9 Click **Save** to keep the created TE explicit path or click **Cancel** to quit without saving.

Delete Explicit Path

Prime Provisioning supports decommission of explicit paths when deleting/decommissioning primary/backup tunnels. This is only supported for IOS XR.

Whether an explicit path can be deleted in such situations depends on whether they are used by other global applications.

Explicit path deletion goes hand in hand with both SR tunnel deletion for primary managed/unmanaged tunnels, backup tunnels, and any non-conformant tunnels and is applicable to all path option types (STRICT, LOOSE, EXCLUDE).

An explicit path configuration will be automatically removed by Prime Provisioning when the explicit path is no longer used by any tunnel in the system due to a change in tunnel configuration. This situation occurs when tunnels are deleted or when tunnels are rerouted in Prime Provisioning.

When the explicit path configuration is removed from the device, the explicit path will still exist in the Prime Provisioning database. Such explicit paths remaining in the database can be reused.

Explicit paths do not get deleted if you reroute or delete the tunnel(s) outside of Prime Provisioning (through CLI on the device itself, for example). However, when a transaction reroutes, deletes, or modifies a tunnel using Prime Provisioning so that an explicit path is no longer used by any tunnels, that explicit path configuration will automatically be removed from the device.

Primary Tunnel Operations

Prime Provisioning allows you to perform a number of primary tunnel operations, which are described in the following sections.

Create Primary Tunnel

After a TE Policy and an explicit path have been set up, a primary tunnel can be created. There are two types of primary tunnels:

- Managed Primary Tunnels
- Unmanaged Primary Tunnels

Below, the GUI flow is described for creating unmanaged primary tunnels. It is very similar for managed primary tunnels and the few differences that exist are described in the section Managed/Unmanaged Primary Tunnels in [Traffic Engineering Management Concepts, page 8-111](#).

To create a managed or an unmanaged primary tunnel, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **Create Managed Tunnel**. The TE Managed Primary Tunnels SR window appears as shown in [Figure 8-14](#).

or

Click **Create Unmanaged Tunnel**. The TE Unmanaged Primary Tunnels SR window appears.

Figure 8-14 Create TE Managed Primary Tunnel

Create TE Managed Primary Tunnel

SR Job ID: New	SR ID: New	SR State: REQUESTED
Creator:	Type: ADD	
Head Device *	Select	
Destination Device *	Select	
Tunnel Policy *	Select	
Tunnel Bandwidth (Kbps):		
Description :		
Tunnel Number:	Auto Gen <input checked="" type="checkbox"/>	
Tunnel ID:		
Customer:		
Auto BW:	Enable: <input type="checkbox"/> Freq (sec): <input type="text"/> Min (Kbps): <input type="text"/> Max (Kbps): <input type="text"/>	
Path Options:		
Showing 1 - 2 of 2 records		
#	Option #	Path Name
1	1	System Path
2	2	Dynamic Path
Rows per page: 10		Page 1 of 1
		Add Delete
		OK Cancel

Note: * - Required Field

The TE Managed Primary Tunnels SR window includes the following elements:

- **Op**—SR operation on the tunnel. This can be one of the following:
 - **ADD**—Indicates a newly added tunnel.
 - **MODIFY**—Indicates a modified existing tunnel.
 - **DELETE**—Indicates an existing tunnel to be deleted.

- **ADMIT**—Indicates an existing tunnel to be admitted by tunnel computation.
- **Tunnel ID**—Unique tunnel identifier used within Prime Provisioning.
- **T#**—Tunnel number on the head router.
- **Head**—Hostname of the head router.
- **Dest**—Hostname of the destination router.
- **Policy**—TE policy for the tunnel.
- **BW**—The tunnel bandwidth. If the tunnel is auto-bw enabled, BW shows the higher of tunnel bandwidth and maximum automatic bandwidth.
- **AutoBW**—Auto Bandwidth enabled if **true**, otherwise **false**.
- **Deploy Status**—Tunnel deployment status.
- **Verified**—Indicates whether tunnel verification was successful (**succeed**, **failed**, or **unknown**).
- **Allow Reroute**—Specifies whether reroute is allowed (**true** or **false**). If reroute is not allowed, the tunnel cannot be set to movable, and hence cannot be rerouted by the operation (placement, grooming, or repair).
- **Head Region**—The region to which the head router belongs.
- **Tail Region**—The region to which the tail router belongs.

The following actions can be performed (buttons):

- **Display**—Open a Topology Display for the network and highlight the selected primary tunnel(s). Selected tunnels are marked in color with directional arrows.
- **Details**—Open the TE Tunnel Details window, which provides type, status, LSP, and other information about the tunnel.
- **Admit**—Admit selected tunnels not previously verified into the managed topology. This feature is used only for discovered tunnels that failed verification or for migrating unmanaged tunnels.
- **Create**—Create a managed primary tunnel.
- **Edit**—Edit a selected primary tunnel.
- **Delete**—Delete selected primary tunnels.
- **Import**—Import tunnel data from import XML file.
- **Placement Tools**—These tools are available only when no change has been made to the tunnels. Apply the following functions against the current topology and tunnels:
 - **Groom**—Analyse the managed tunnels in the network and reroute them to reduce the maximum link utilization.
 - **Tunnel Audit**—Determine if changes to previously made SRLGs or backup tunnels have caused constraint violations in managed tunnels (this can occur when managed tunnels have FRR protection constraints).
 - **Tunnel Repair**—Repair any managed tunnel constraint violations revealed by **Placement Tools > Tunnel Audit**.
- **Update Tunnel ID**—Update Tunnel ID(s) directly in the repository without deploying the corresponding tunnel(s).
- **Proceed with Changes**—For verifying changes in tunnels. When tunnels have been created, deleted, admitted, or their attributes altered, you can proceed with one of the following placement tools:
 - **Tunnel Audit**—Checks what constraint violations modifications to tunnels might cause.

- **Tunnel Placement**—Admit new tunnels and modify tunnels already admitted into the network.
- **Tunnel Repair**—Resolve inconsistencies caused by changes to bandwidth requirements or delay parameters of existing tunnels by moving as few existing tunnels as possible to accommodate the changes.

Note that for the unmanaged tunnels list, the last two columns in the managed tunnels list (Verified and Allow Reroute) are replaced by the Conformance column.

In the following example, an unmanaged tunnel is created.

Step 3 Click **Create**.

The Create TE Unmanaged Primary Tunnel window appears.

The Create TE Managed Primary Tunnel window and Create TE Unmanaged Primary Tunnel window have only minor differences and include the following elements:

- **Head Device**—Head device for the tunnel.
- **Destination Device**—Destination device for the tunnel.
- **Tunnel Policy**—A set of rules established for a tunnel.
- **Tunnel Bandwidth**—Total allocated bandwidth of the tunnel.
- **Description**—Descriptive text to help identify the tunnel.
- **Tunnel Number**—Tunnel number corresponding to the tunnel interface name.
 - **Auto Gen**—Check this box to generate the tunnel number automatically. Otherwise, enter a desired number.



Note If a manually entered tunnel number is too low, it could prevent deployment.



Note MPLS-TE tunnels can potentially interfere with multicast GRE tunnels. Prime Provisioning creates new tunnels using auto-gen and this tunnel number might already be used by an MDT GRE tunnel. As a result, Prime Provisioning uses high tunnel numbers to avoid any complications.

- **Tunnel ID**—Unique tunnel identifier used within Prime Provisioning.
- **Customer**—Selected customer for the tunnel.
- **Auto BW**—A way to configure a tunnel for automatic bandwidth adjustment and to control the manner in which the bandwidth for a tunnel is adjusted.
 - **Enable**—Check this box to enable automatic bandwidth.
 - **Freq**—Interval between bandwidth adjustments.
 - **Min**—Minimum automatic bandwidth, in kbps, for this tunnel.
 - **Max**—Maximum automatic bandwidth, in kbps, for this tunnel.

Path options:

- **Option #**—Sequential number of available explicit paths.
- **Path Name**—Name of the explicit path. In case of an existing path, the name is a URL that links to the Explicit Path Viewer.

- **System Path**—System generated explicit path. For managed tunnels, the first path has to be an explicit path. If a tunnel contains a system path, the planning function will generate an optimal path for the tunnel.
- **Dynamic Path**—A dynamic path is provisioned by allowing the head router to find a path. The **dynamic** keyword is provisioned to the routers.
- **Path Type**—Path option type, Explicit or Dynamic.
- **Lock Down**—Disables reoptimization check on the tunnel, if checked, meaning the path cannot be changed.

Step 4 To select a **Head Device** in the Create TE Unmanaged Primary Tunnel window, click the corresponding **Select** button to open the Select Device for TE Head Router window.

Step 5 Select a device name and click **Select**.

The Select Device for TE Head Router window closes and the prompt returns to the Create TE Unmanaged Primary Tunnel window.

Step 6 To select a **Destination Device** in the Create TE Unmanaged Primary Tunnel window, click the corresponding **Select** button to open the Select Device for TE Tail Router window.

Step 7 Select a device name and click **Select**.

The Select Device for TE Tail Router window closes and the prompt returns to the Create TE Unmanaged Primary Tunnel window.

Step 8 To select a **Tunnel Policy** in the Create TE Unmanaged Primary Tunnel window, click the corresponding **Select** button to open the Select Unmanaged TE Tunnel Policy window.



Note

When creating a managed tunnel, make sure that one or more managed tunnel policies are available. If that is not the case, go to **Policies** (see [Create TE Policy, page 8-28](#)) and make sure to check the **Managed** check box.

Step 9 Select a policy and click the **Select** button.

This brings you back to the tunnel editor.

Step 10 Click **Add** to set up path options for the tunnel. The Select TE Explicit Path window appears.

The **Path Options** section provides two path types:

Explicit Path—A fixed path from a specific head to a specific destination device that includes three types of paths: **Strict**, **Loose**, and **Exclude**.

Dynamic Path—A dynamic path is provisioned by allowing the head router to find a path. The **dynamic** keyword is provisioned to the routers.

Step 11 Select the desired TE Explicit Path unless you prefer dynamic path only.

If none is available, you can set one up first. To do so, see [Create Explicit Path, page 8-29](#).

Step 12 Click **Select**.

The selected path appears in the **Path Options** section of the create window.

For explicit paths (<head_device>-<destination_device>), you can click the pathname to open the non-editable Explicit Path Viewer.

For an explanation of the various window elements, see [Create Explicit Path, page 8-29](#).

Step 13 In the Create TE Unmanaged Tunnel window, click **OK** to accept the entered tunnel information or click **Cancel** to quit and return to the TE Unmanaged Primary Tunnels SR window.

The TE Unmanaged Primary Tunnel SR window appears with the newly created SR with the **Op** field set to **ADD**.



Note The added tunnel can be reverted from the **ADD** state to its original state by selecting it and clicking **Delete**. The tunnel is removed from the tunnel list.

Step 14 In the TE Unmanaged Primary Tunnel window, click **Save & Deploy** (see [Note](#) on page 36) to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more primary tunnels and then save and deploy all changes.

When you click **Save & Deploy**, Prime Provisioning locks the TE routers effected, which will block any subsequent SRs which use that TE Router until the SRs are finished. It is safe to try and deploy other SRs in the system. If there is any conflict with the SR currently being processed, Prime Provisioning will simply ask you to wait until it is complete.

To see the state of deployment, go to the Service Requests window at **Operate > Service Request Manager** or open **Operate > Task Manager**.

- **Save & Deploy**—For committing tunnel changes that do not impact tunnel placement. There are two options for saving and deploying SR tunnels to the network:
 - **SR Tunnels Only**—Deploy all tunnel changes that does not impact tunnel placement, or if no changes were made to the SR, use this to redeploy the SR that was in **Requested** or **Invalid** state.
 - **Force Deploy All Tunnels**—Force deployment of all tunnels in this SR. This could be useful when previous provisioning of the SR has failed, so that it is necessary to force through the deployment of all tunnels in the SR.



Note You might see Elixir Warnings during TE Tunnel deployment. The deployment will be successful and the warning messages can safely be ignored.



Note For managed tunnels, you cannot deploy the service request until you have used the **Proceed with Changes** button to perform either Tunnel Placement, Tunnel Audit, or Tunnel Repair (see [Advanced Primary Tunnel Management, page 8-44](#)).



Note With the exception of TE Traffic Admission SRs, TE SRs are always deployed immediately from the specific TE SR window, not from **Operate > Service Request Manager**.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR (first **REQUESTED**, then **PENDING**, then **DEPLOYED**, if successful).

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

If the SR does not go to the **Deployed** state, go to the Task Logs window to see the deployment log (**Operate > Task Manager > Logs**) as described in [SR Deployment Logs, page 11-47](#).

To edit the service request from the **Service Request Manager** window, go back to the TE Managed Primary Tunnels SR or the TE Unmanaged Primary Tunnels SR window as described in [Edit Primary Tunnel, page 8-37](#).

Edit Primary Tunnel

Primary tunnel attributes can be modified in the primary tunnel editor.

There are two ways to access the primary tunnel editor:

- from the managed or unmanaged primary tunnels SR window or
- from the Service Requests window.

Access from Primary Tunnel SR Window

To access the primary tunnel editor from the primary tunnel SR window (TE Managed Primary Tunnels SR or TE Unmanaged Primary Tunnels SR window) and edit a managed or an unmanaged primary tunnel, use the following steps:

-
- Step 1** Choose **Traffic Engineering**.
- Step 2** Click **Create Managed TE Tunnel**. The TE Managed Primary Tunnels SR window in [Figure 8-14](#) appears.
- or
- Click **Create Unmanaged TE Tunnel**. The TE Unmanaged Primary Tunnels SR window appears.
- Step 3** To edit a tunnel SR, select the desired SR and click **Edit**.
- The Edit TE Managed Primary Tunnel or the Edit TE Unmanaged Primary Tunnel window appears.
- The primary tunnel editor is identical to that of the create primary tunnel GUI. For an explanation of the various window elements, see [Create Primary Tunnel, page 8-32](#).
- Step 4** Make the desired changes and click **OK** to accept, or **Cancel** to discard the changes.
- In the TE Unmanaged Primary Tunnel SR window, the **Op** field changes to MODIFY.



Note The modified tunnel can be reverted to its original state by selecting it and clicking **Delete**. The MODIFY flag in the Op column disappears.

- Step 5** Click **Save & Deploy** to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more primary tunnels and then save and deploy all changes.
- The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.
- For more information on working with service requests, see the managing service requests part elsewhere in this guide.
-

Access from Service Requests Window

To access the primary tunnel editor from the Service Requests window, assuming that the SR has been created, use the following steps:

-
- Step 1** Choose **Operate > Service Request Manager**.
- Step 2** To edit the desired tunnel SR, select the SR in question and click **Edit**.
Depending on whether a managed or an unmanaged tunnel has been selected, the TE Managed Primary Tunnel SR or the TE Unmanaged Primary Tunnel SR window appears displaying the SR selected in the Service Requests window.
- Step 3** Select the tunnel SR and click **Edit**.
The Edit TE Unmanaged Primary Tunnel window appears.
Go to [Access from Primary Tunnel SR Window, page 8-37](#) and continue the process from [Step 4](#).
-

Delete Primary Tunnel

TE tunnels can be deleted either from the TE Links List window (see [Deleting TE Tunnels, page 8-25](#)) or in the primary or backup tunnels SR windows.

To delete a managed or an unmanaged primary tunnel from the TE Managed Primary Tunnels SR or TE Unmanaged Primary Tunnels SR window, use the following steps:

-
- Step 1** Choose **Traffic Engineering**.
- Step 2** Click **Create Managed TE Tunnel**. The TE Managed Primary Tunnels SR window appears.
or
Click **Create Unmanaged TE Tunnel**. The TE Unmanaged Primary Tunnels SR window appears.
- Step 3** To delete a tunnel, select the desired tunnel(s) and click **Delete**.
The **Op** field status changes to **DELETE**.
For an explanation of the various window elements, see [Create Primary Tunnel, page 8-32](#).



Note The deleted tunnel can be reverted to its original state by selecting it and clicking **Delete**. The DELETE flag in the Op column disappears.

- Step 4** Click **Save & Deploy** to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more primary tunnels and then save and deploy all changes.
The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.
For more information on working with service requests, see the managing service requests part elsewhere in this guide.
-

Backup Tunnel Operations

Prime Provisioning allows you to perform a number of backup tunnel operations, which are described in this section.

[Traffic Engineering Management Concepts, page 8-111](#) contains a section on Connectivity Protection (CSPF) Backup Tunnels, which is one of the techniques used to provide backup protection.

Create Backup Tunnel

Backup tunnels are created in much the same way as primary tunnels. In both cases, building an explicit path is not required when an existing path already traverses the desired routers. A path can be used for any number of tunnels within its bandwidth capacity.

A precondition for creating a backup tunnel is the presence of an explicit path. To create an explicit path, see [Create Explicit Path, page 8-29](#).

To create a backup tunnel, use the following steps:

Step 1 Choose **Traffic Engineering > Create TE Backup Tunnel**.

The TE Protection SR window appears.

The TE Protection SR window includes the following elements:

The columns in the tunnel list provides the following information:

- **Op**—Current SR operation on the tunnel. This can be one of the following:
 - **ADD**—Indicates a newly added tunnel, either calculated by the system or entered by the user.
 - **MODIFY**—Indicates a modified existing tunnel.
 - **DELETE**—Indicates an existing tunnel to be deleted, either computed by the system or originated by the user.
- **Tunnel ID**—Unique tunnel identifier used within Prime Provisioning.
- **T#**—Tunnel number on the head router.
- **Head**—Hostname of the head router.
- **Dest**—Hostname of the destination router.
- **BW Quota**—Amount of bandwidth that this backup tunnel can protect. The router limits the LSPs that can use this backup tunnel so that the sum of the bandwidth of the LSPs does not exceed the specified amount of bandwidth. If there are multiple backup tunnels, the router will use the best-fit algorithm.
- **Deploy Status**—Tunnel deployment status.
- **Conformance**—Indicates whether the tunnel is found to be conformant when running discovery. A tunnel is non-conformant if it has a non-zero bandwidth reservation and a zero hold or setup priority. If a tunnel is entered through TEM, it is always conformant. A connectivity protection tunnel is marked Conformance = true if it has zero tunnel bandwidth, unlimited backup bandwidth, and an 'exclude address' first path option. Otherwise, it is marked Conformance = false.
- **Backup Type**—Can be either bandwidth protected backup tunnels (**BW Protected**) or CSPF-routed backup tunnels (**CSPF**). For more information about these types of backup tunnels, see [Traffic Engineering Management Concepts, page 8-111](#).
- **Head Region**—The region to which the head router belongs.

- **Tail Region**—The region to which the tail router belongs.

Step 2 Click **Create**.

The Create TE Backup Tunnel window in [Figure 8-15](#) appears.

Figure 8-15 Create TE Backup Tunnel

Create TE Managed Primary Tunnel

SR Job ID: New		SR ID: New		SR State: REQUESTED	
Creator:		Type: ADD			
Head Device *	<input type="button" value="Select"/>				
Destination Device *	<input type="button" value="Select"/>				
Tunnel Policy *	<input type="button" value="Select"/>				
Tunnel Bandwidth (Kbps):	<input type="text"/>				
Description:	<input type="text"/>				
Tunnel Number:	Auto Gen <input checked="" type="checkbox"/>	<input type="text"/>			
Tunnel ID:	<input type="text"/>				
Customer:	<input type="text"/>				
Auto BW:	Enable: <input type="checkbox"/>				
	Freq (sec):	<input type="text"/>			
	Min (Kbps):	<input type="text"/>			
	Max (Kbps):	<input type="text"/>			
Path Options:					
Showing 1 - 2 of 2 records					
#	<input type="checkbox"/> Option #	Path Name	Path Type	Lock Down	
1	<input type="checkbox"/> 1	System Path	Explicit	<input type="checkbox"/>	
2	<input type="checkbox"/> 2	Dynamic Path	Dynamic	<input type="checkbox"/>	
Rows per page: 10		<input type="button" value="Previous"/> <input type="button" value="Next"/> Page 1 of 1			
<input type="button" value="Add"/> <input type="button" value="Delete"/>					
<input type="button" value="OK"/> <input type="button" value="Cancel"/>					

Note: * - Required Field

The Create TE Backup Tunnel window includes the following elements:

- **Head Device**—Head device for the tunnel.
- **Destination Device**—Destination device for the tunnel. The selection window is very similar to the Head Device selection window.
- **Protected Interface(s)**—Interface(s) on the head router that this backup tunnel protects.
- **Description**—Descriptive text to help identify the tunnel.
- **Backup Bandwidth Limit**—Bandwidth protected by the backup tunnel.
 - **Any Pool BW**—Bandwidth set aside for the protection of either the Sub Pool or the Global Pool.
 - **Sub Pool (BC1) BW**—Bandwidth set aside for the Sub Pool.
 - **Global Pool (BC0) BW**—Bandwidth set aside for the Global Pool.

For a definition of pool types, see [Traffic Engineering Management Concepts, page 8-111](#).

- **Tunnel Number**—Tunnel number corresponding to the tunnel interface name.
 - **Auto Gen**—Check this box to generate the tunnel number at provisioning time. Otherwise, enter a desired number.



Note If a manually entered tunnel number is too low, it could prevent deployment.

- **Tunnel ID**—Unique tunnel identifier used within Prime Provisioning.
- **Tunnel Bandwidth**—Total allocated bandwidth of this backup tunnel (display only).
- **Tunnel Pool Type**—Tunnel bandwidth pool type for this policy (display only). For a definition of pool types, see [Traffic Engineering Management Concepts, page 8-111](#).
 - **Global Pool (BC0)**—Bandwidth will be reserved from Global Pool.
 - **Sub Pool (BC1)**—Bandwidth will be reserved from Sub Pool.
- **Setup Priority (0-7), Hold Priority (0-7), Affinity, Affinity Mask**—All manually created backup tunnels should have setup and hold priorities of 0 and affinity value and mask of 0x0 for them to be able to protect an element.

Path options:

- **Option #**—Sequential number of available explicit paths.
- **Path Name**—Name of the explicit path.
- **Path Type**—Explicit path type (**Explicit** or **Dynamic**)
- **Lock Down**—Disables reoptimization check on the tunnel, if checked.

Step 3 Select, at a minimum, a **Head Device**, a **Destination Device**, and a **Protected Interface**.

Also, specify a **Backup Bandwidth Limit** greater than zero. Add other tunnel information as desired.

Step 4 Click **Add** to add just one path.

The Select TE Explicit Path window appears.

Step 5 Select an explicit path.

It must match the head and destination of an existing path. If none is available, you first must set one up. To do so, see [Create Explicit Path, page 8-29](#).

Step 6 Click **Select**.

The selected path appears in the **Path Options** section of the page as shown in the Select TE Explicit Path window.

For explicit paths, you can click the pathname to open the Explicit Path Viewer.

Step 7 In the Create TE Backup Tunnel window, click **OK** to accept the entered tunnel information or click **Cancel** to quit the window without saving it.

In the TE Protection SR window, a new backup tunnel is added in the tunnel list with the **Op** field set to ADD.



Note

The added tunnel can be reverted to its original state by selecting it and clicking **Delete**. The tunnel is removed from the tunnel list.

Step 8 Click **Save & Deploy** to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more backup tunnels and then save and deploy all changes.

The **Save & Deploy** button provides two options:

- **SR Tunnels Only**—Deploy all tunnel changes that does not impact tunnel placement, or if no changes were made to the SR, use this to redeploy the SR that was in **Requested** or **Invalid** state.
- **Force Deploy All Tunnels**—Force deployment of all tunnels in this SR. This could be useful when previous provisioning of the SR has failed, so that it is necessary to force through the deployment of all tunnels in the SR.

When you click **Save & Deploy**, Prime Provisioning locks the TE routers effected, which will block any subsequent SRs which use that TE router until the SRs are finished. It is safe to try and deploy other SRs in the system. If there is any conflict with the SR currently being processed, Prime Provisioning will simply ask you to wait until it is complete. To see the state of deployment, go to the Service Requests window under Inventory and Connection Manager or open the Task Manager under Monitoring.

**Note**

You might see Elixir Warnings during TE Tunnel deployment. The deployment will be successful and the warning messages can safely be ignored.

**Note**

With the exception of TE Traffic Admission SRs, TE SRs are always deployed immediately from the specific TE SR window, not from the **Operate > Service Request Manager** page.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

If the SR does not go to the **Deployed** state, go to the Task Logs window to see the deployment log (**Operate > Task Manager > Logs**) as described in [SR Deployment Logs, page 11-47](#).

Edit Backup Tunnel

Backup tunnel attributes can be modified in the backup tunnel editor.

There are two ways to access the backup tunnel editor:

- from the Protection SR window or
- from the Service Requests window.

From the Protection SR Window

To access the Protection SR window to edit a backup tunnel, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Create TE Backup Tunnel**.
The TE Protection SR window appears.
- Step 2** To edit a tunnel SR, select the desired SR and click **Edit**.
The Edit TE Backup Tunnel window appears. The backup tunnel editor is identical to that of the create backup tunnel GUI. For an explanation of the various window elements, see [Create Backup Tunnel, page 8-39](#).
- Step 3** Make the desired changes and click **OK**.
In the TE Protection window, the **Op** field changes to MODIFY.

**Note**

The modified tunnel can be reverted to its original state by selecting it and clicking **Delete**. The MODIFY flag in the Op column disappears.

- Step 4** In the TE Protection SR window, click **Save & Deploy** to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more backup tunnels and then save and deploy all changes.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

From the Service Requests Window

To edit a backup tunnel from the **Service Requests** window, assuming that the SR has been created use the following steps:

-
- Step 1** Choose **Operate > Service Request Manager**.
- Step 2** To edit the desired tunnel SR, select the SR in question and click **Edit**.
- The TE Protection SR window appears displaying the SR selected in the Service Request Manager window.
- Step 3** Select the tunnel SR and click **Edit**.
- The Edit TE Backup Tunnel window appears.
- Go to [Edit Backup Tunnel, page 8-42](#) and continue the process from [Step 3](#).
-

Delete Backup Tunnel

TE tunnels can be deleted either from the TE Links List window (see [Deleting TE Tunnels, page 8-25](#)) or in the primary or backup tunnels SR windows.

To delete a backup tunnel from the TE Protection SR window, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Create TE Backup Tunnel**.
- The TE Protection SR window appears.
- Step 2** To delete a tunnel SR, select the desired SR and click **Delete**.
- The **Op** field status changes to **DELETE** for unmanaged tunnels.
- For an explanation of the various window elements, see [Create Backup Tunnel, page 8-39](#).



Note The deleted tunnel can be reverted to its original state by selecting it and clicking **Delete**. The DELETE flag in the Op column disappears.

Click **Save & Deploy** to either deploy the new tunnel SR to the network or force deploy all tunnels, or you can create or edit more primary tunnels and then save and deploy all changes.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Purging a Service Request

The Purge operation in the Service Request Manager window is designed to remove a service request from the repository without affecting the network.

The **Purge** button has 2 options:

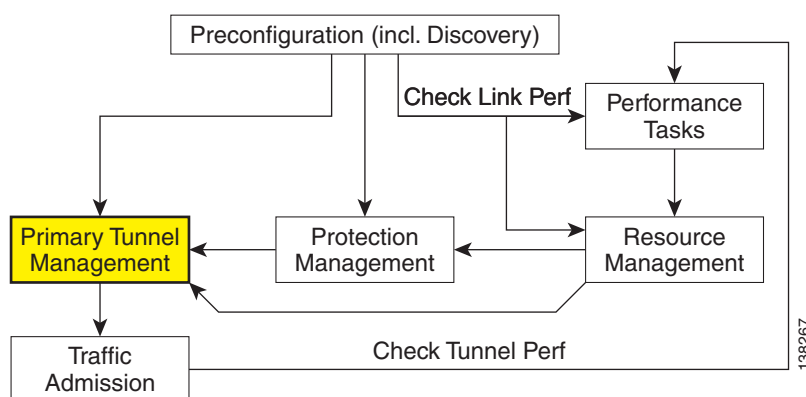
- **Purge**—The regular purge can only be used on the service request in **CLOSED** state. Therefore, it cannot be used on TE Resource, TE Tunnel, or TE Protection service requests because these cannot be decommissioned. These three types of service requests can only be force purged.
- **Force Purge**—During force purge, the repository checks the necessary dependency on the service request before it can be purged, so if a service request cannot be purged, there will be an error message.

Advanced Primary Tunnel Management

In addition to the basic tunnel management tools described in [Basic Tunnel Management, page 8-27](#), Prime Provisioning gives access to a set of advanced tunnel planning tools that provide optimal placement of tunnels to ensure efficient use of network resources.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning primary tunnel management occurs.

Figure 8-16 Prime Provisioning Process Diagram - Primary Tunnel Management



The advanced tools are available for managed tunnels only. The difference between managed and unmanaged tunnels is described in the section Managed/Unmanaged Primary Tunnels in [Traffic Engineering Management Concepts, page 8-111](#).

This section includes the following:

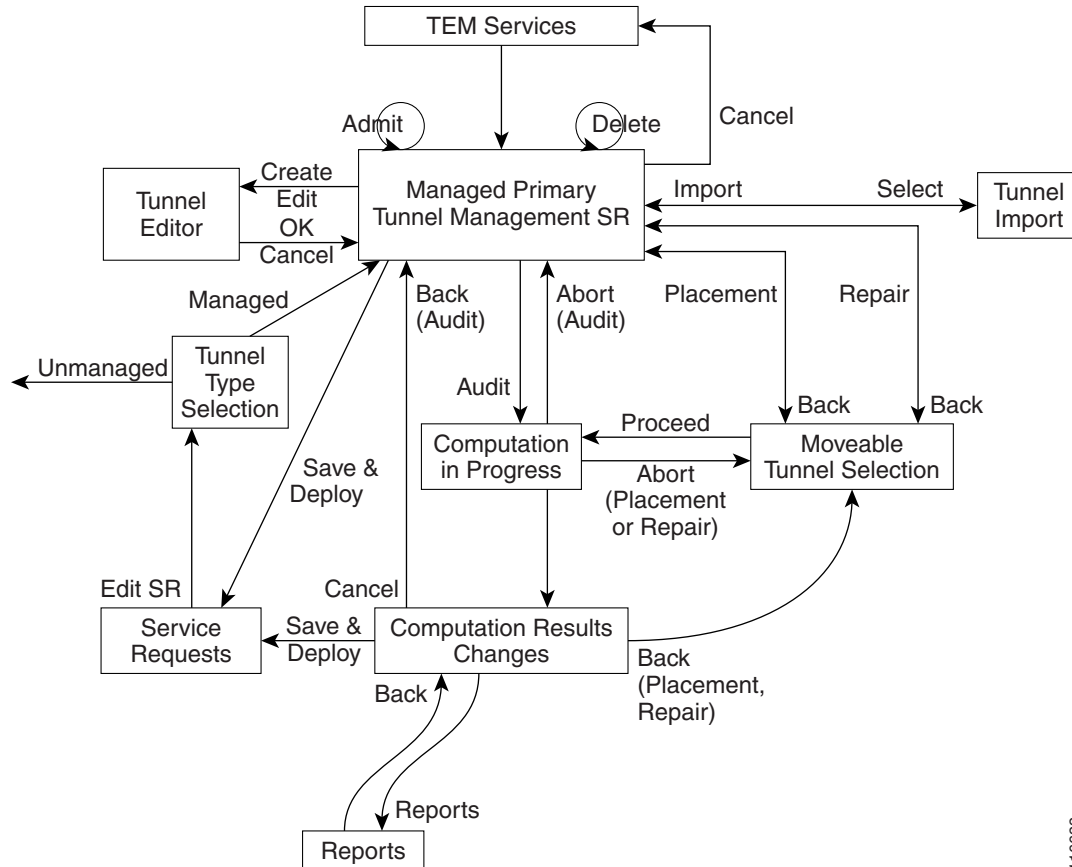
- [Tunnel Operations, page 8-45](#)
 - [Create Primary Tunnel, page 8-46](#)

- Edit Primary Tunnel, page 8-49
 - Delete Primary Tunnel, page 8-49
 - Admit Primary Tunnel, page 8-49
 - Import Primary Tunnel, page 8-49
- Planning Strategy, page 8-51
- Placement Tools, page 8-52
 - Tunnel Audit, page 8-52
 - Tunnel Placement, page 8-55
 - Tunnel Repair, page 8-56
 - Grooming, page 8-57.

Tunnel Operations

This section explains the advanced tunnel operations in Prime Provisioning that incorporate the planning tools.

An overview of the primary tunnel management process is provided in [Figure 8-17](#).

Figure 8-17 Primary Tunnel Management Processes

116622

For **Tunnel Type Selection**, when you select **Unmanaged** the TE Unmanaged Primary Tunnel SR window appears (see [Basic Tunnel Management](#), page 8-27).

All other elements in [Figure 8-17](#) are described in this section.

Create Primary Tunnel

To create a TE managed primary tunnel with the RG license installed, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **Create Managed TE Tunnel**.

The TE Managed Primary Tunnels SR window appears.

For an explanation of the various window elements, see [Create Primary Tunnel](#), page 8-32.

Step 3 Click **Create**.

The Create TE Managed Primary Tunnel window appears.

For an explanation of the various window elements, see [Create Primary Tunnel](#), page 8-32.

The **Path Options** section provides three path types, **System Path**, **Explicit Path**, and **Dynamic Path**.

A **System Path** is an Prime Provisioning system generated explicit path (immovable). The first path has to be an explicit path.

An **Explicit Path** is a fixed path from a specific head to a specific destination device.

A **Dynamic Path** is provisioned by allowing the head router to find a path. The **dynamic** keyword is provisioned to the routers.

Step 4 To select a **Head Device**, click the corresponding **Select** button to open the device selection window. Select a head device and click **Select**.

Step 5 To select a **Destination Device**, click the corresponding **Select** button to open the device selection window.

Select a tail device and click **Select**.

Step 6 To select a **Tunnel Policy**, click the corresponding **Select** button to open the policy selection window.



Note

If no tunnel policies are available, the reason could be that they are all unmanaged. To create a managed tunnel, first create a managed policy in **Service Design > Policy Manager** (see [Create Policy, page 8-71](#)) by making sure to check the **Managed** check box.

The Select Managed TE Tunnel Policy window includes the following elements:

- **Policy Name**—Name of the TE policy.
- **Pool Type**—Tunnel bandwidth pool type for this policy. For a definition of pool types, see the Bandwidth Pools section in [Traffic Engineering Management Concepts, page 8-111](#).
 - **SUB_POOL**—Bandwidth will be reserved from Sub Pool.
 - **GLOBAL**—Bandwidth will be reserved from Global Pool.
- **Setup Priority**—Priority used when signaling an LSP for the tunnel to determine, which of the existing tunnels can be preempted. Valid values are from 0 to 7, where a lower number indicates a higher priority. Therefore, an LSP with a setup priority of 0 can preempt any LSP with a non-0 hold priority.
- **Hold Priority**—Priority associated with an LSP for the tunnel to determine if it should be preempted by other LSPs that are being signaled. Valid values are from 0 to 7, where a lower number indicates a higher priority.
- **Affinity**—Attribute values required for links carrying the tunnel (bit values are either 0 or 1).
- **Affinity Mask**—Attribute values to be checked. If a bit in the mask is 0, a link's attribute value of that bit is irrelevant. If a bit in the mask is 1, the link's attribute value and the tunnel's required affinity for that bit must match.
- **Delayed Constraint**—True or false value. If true, the tunnel has a maximum delay that its path must not exceed.
- **FRR Protection**—Used to enable an MPLS traffic engineering tunnel to use a backup tunnel in the event of a link failure if a backup tunnel exists.
 - **None**—No backup tunnel needed.
 - **Best Effort**—Use backup tunnel if available.
 - **Link and SRLG (only managed tunnels)**—Specifies that primary tunnels should be routed only through links and SRLGs that are protected by FRR backup tunnels.
 - **Link, SRLG and Node (only managed tunnels)**—Specifies that primary tunnels should be routed only through links, SRLGs and nodes that are protected by FRR backup tunnels.

- **MPLS IP Enabled**—Specifies whether MPLS IP has been configured for the corresponding tunnel.

Step 7 Specify a tunnel bandwidth greater than zero.

Step 8 Add other tunnel information as desired.

Step 9 Optionally, if you want to specify an explicit path rather than using the system path provided by Prime Provisioning, delete the system path and subsequently add the explicit path.

For a more detailed explanation of this step, see [Create Primary Tunnel, page 8-32](#).

Step 10 In the Create TE Managed Tunnel window, click **OK** to accept the entered tunnel information or **Cancel** to quit and return to the TE Managed Primary Tunnels SR window.

The TE Managed Primary Tunnel SR window appears displaying the new tunnel with the **Op** field set to ADD to signify that an SR has been added.



Note The added tunnel can be reverted to its original state by selecting it and clicking **Delete**. The tunnel is removed from the tunnel list.

Step 11 In the TE Managed Primary Tunnel SR window, you can create or edit more tunnels, or if you are done with all the changes, proceed in one of the following two ways depending on which of the following buttons are active (**Save & Deploy** is not available after the **Create** operation):

- **Proceed with Changes:** The changes you entered impacts tunnel placement. Click on this to continue with one of the planning flows described in the Placement Tools (see [Placement Tools, page 8-52](#)) until the SR can be saved and deployed.
- **Save & Deploy:** The changes you entered do not impact tunnel placement. Click on this to save and deploy the SR. This function is further described in [Create Primary Tunnel, page 8-32](#).

When you click **Save & Deploy**, Prime Provisioning locks the TE routers effected, which will block any subsequent SRs which use that TE router until the SRs are finished. It is safe to try and deploy other SRs in the system. If there is any conflict with the SR currently being processed, Prime Provisioning will simply ask you to wait until it is complete. To see the state of deployment, go to the Service Requests window under Inventory and Connection Manager or open the Task Manager under Monitoring.



Note With the exception of TE Traffic Admission SRs, TE SRs are always deployed immediately from the specific TE SR window, not from the Service Requests page in **Inventory and Connection Manager**.

If **Save & Deploy** was selected in [Step 11](#), the Service Requests window (**Operate > Service Request Manager**) opens and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.



Note You might see Elixir Warnings during TE Tunnel deployment. The deployment will be successful and the warning messages can safely be ignored.

If the SR does not go to the **Deployed** state, go to the Task Logs window to see the deployment log (**Operate > Task Manager > Logs**) as described in [Task Logs, page 11-28](#).

Edit Primary Tunnel

The only difference between creating and editing tunnels is that in the tunnel editor, the head and destination devices and tunnel number fields are not editable. Otherwise, you create and edit the same attributes.

Only **Proceed with Changes** or **Save & Deploy**, not both, are available depending on whether the changes you entered impacts tunnel placement.

To edit a primary tunnel, see [Edit Primary Tunnel, page 8-37](#)

Delete Primary Tunnel

To delete one or more tunnels, see [Delete Primary Tunnel, page 8-38](#).

Admit Primary Tunnel

The Admit function is used to admit selected tunnels not previously verified into the managed topology. This feature is used only for discovered tunnels that failed verification. During the discovery process, verification is performed with the Tunnel Placement algorithm, as if the tunnels were admitted for the first time.

Verification means that the discovered managed tunnel is verified against the network topology and TEM checks if there is enough bandwidth along the tunnel path (both are specified in the tunnel).

In general, verification will fail if there is not enough bandwidth due to the existence of other tunnels or a limitation on link capacity/bandwidth.

More specifically, this can happen when a priority 0 tunnel is created independently of TEM and a TE Discovery task is run. If the tunnel does not satisfy all the managed tunnel constraints (that is, if it is reserving more bandwidth than is available in a link that it passes through) TE discovery will mark it as 'verified = false'. It will not be managed by TEM until you use the Admit button to make it verified. Typically this would have to be accompanied with some other tunnel or resource change to ensure that the constraint is now satisfied.

To admit a primary tunnel, use the following steps:

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | In the TE Managed Primary Tunnel SR , select one or more unverified tunnels to migrate. |
| Step 2 | Click Admit .

The unverified tunnel(s) are verified and, if successful, an ADMIT flag will appear in the Op column. |
| Step 3 | Choose Proceed with Changes > Tunnel Placement to determine if the tunnels can be placed. If not, edit the tunnels and try again. |
-

Import Primary Tunnel

This feature allows you to update tunnels in bulk through a file-based import mechanism. The data is migrated into the managed primary tunnel service request.

Construct XML Import File

To import tunnels from a file, first construct an XML import file conforming to the structure defined in the system supplied Document Type Definition (DTD) file (see [Document Type Definition \(DTD\) File, page 8-108](#)), and save the XML file together with the DTD file on the Prime Provisioning server under the same directory. To create a valid import file, use the provided command line validation tool (see [Command Line Validation Tool, page 8-50](#)).

The following files are necessary for importing data into the Prime Provisioning application and are included in the installation:

- DTD file for the import file in
`<installedDir>/resources/java/xml/com/cisco/vpnsd/ui/te`
 - **TeImport.dtd**
(a sample file, 'sample.xml', is also included)
- Shell script for executing the command line validator in the `<installedDir>/bin` directory.
 - **ImportTeTunnels**
Usage: **importTeTunnels** *<importfile>*

importfile is a XML file and must specify **TeImport.dtd** as its DTD. **TeImport.dtd** must be in the same directory as *importfile*.

Command Line Validation Tool

The purpose of a command line validator is to help construct a valid import file off-line that corresponds to **TeImport.dtd**. The tool helps screen out errors associated with files that are not well-formed and files that do not conform to the rules set by the DTD.

For instructions on how to use the DTD file, see the DTD file documentation.

The tool reads the import file line-by-line, echoes each line in on the output as it parses, and reports any parsing error it encounters. The parsing and validation continues even when parsing errors are encountered for as long as the file structure makes sense.



Note

This tool does not check for cross field validation or data integrity errors with respect to the Prime Provisioning application.

Import Procedure

The file-based import feature is only enabled when there are no uncommitted new, changed, or deleted tunnels in the service request.

It provides a way of adding, editing, deleting, or migrating many tunnels at a time.

To start the import procedure, use the following steps:

-
- Step 1** Prepare the XML import file in accordance with the DTD file.
 - Step 2** Go to **Traffic Engineering**.
 - Step 3** Select provider if this has not been done earlier in the session.
 - Step 4** Click **Create Managed TE Tunnel**.

The TE Managed Primary Tunnels SR window appears.

Step 5 Click **Import** to start the import process.

The Select Import File window appears.



Note The Import button is only enabled when there are no uncommitted new, changed, or deleted tunnels in the service request.

The Select Import File window lists all the XML files and any directories under the directory name shown in the **Look in** field.

The default directory shown in the **Look in** field corresponds to the installation directory in which the DTD and sample XML files reside.

Step 6 Select the desired XML file to be used for the import operation.

The system then parses the file. If any error is detected, it will be reported in the Tunnel Import Error Status window.

The Tunnel Import Error Status window shows the URL of the file, its last modified timestamp, the import status, and any error/warning messages.

Step 7 If the import operation failed, click **Cancel** to return to the previous window.

If it is partially successful, the **Continue** button is enabled, thereby providing an additional option to accept system treatment for errors/warnings and continue with the import operation.

Step 8 If the file is parsed successfully or you click **Continue**, all valid tunnels in the file are added to the service request and the TE Managed Primary Tunnels SR window is re-displayed in the SR view. The imported tunnels are displayed with the appropriate tunnel **Op** type.

Planning Strategy

The main objective of using the planning tools is to achieve optimal overall network utilization while causing minimal impact on any existing traffic on the network.

In most cases, the following strategy can be applied:

- Attempt to admit the new traffic optimizing on utilization (Placement feature) without allowing existing traffic to be moved. This offers the possibility of accommodating the new traffic without any changes to the existing traffic, while still optimising reserved bandwidth utilization under the constraint that existing tunnels do not move.
- If this fails, attempt to admit the same new traffic minimizing change to existing traffic (Repair feature) to see if the new traffic can be accommodated without affecting any more existing tunnels than necessary.
- If this succeeds in placing the new traffic, but you feel that the overall reserved bandwidth utilization is higher than would be preferred, consider grooming the network.
- If the Repair fails, review the parameters that control how many changes can be considered. Alternatively the specification to the desired traffic could be changed, or resource modifications could be made.

This strategy reflects the different approaches taken by the different algorithms in searching for solutions. However, other combinations are possible.

Placement Tools

Planning tools for primary tunnels are available from the **Proceed with Changes** and **Placement Tools** buttons in the TE Primary Tunnel SR window depending on whether a change has been made to the managed primary tunnels.

- **Proceed with Changes:** Used when you have made changes (add/change/delete/admit) to the tunnels. Tunnel operations are described in [Tunnel Operations, page 8-45](#). Then choose one of the placement tools to verify primary placement with the system and continue with deployment. This button is also available in Resource Management.
- **Placement Tools:** Used to perform planning functions on the existing network.
 - The Tunnel Audit option should be used to verify the constraint-based placement of existing managed primary tunnels with the existing network topology. You can use this option to find out the optimality of your primary placement. If you are requiring protection levels above "Best Effort" on your primary tunnels, it is also important to perform an audit after any changes have been made in the protection network. If the audit results in warnings/violations, you can use the Tunnel Repair option help you find a solution.
 - The Groom option is used for optimizing your primary placement. In all primary computation, a quality report is produced which displays the optimality and utilization of the bandwidth pools. You can perform a Tunnel Audit first to determine if grooming is needed on your network.

The planning tools are described in detail in the following sections.

**Note**

If tunnel attributes that are not supported by the placement tools (such as auto-bw frequency) are changed in conjunction with attributes that are supported, the attributes appear correctly in the TE Computation Results window. But if only unsupported attributes are changed, the TE Computation Results window still shows no achieved changes and the **Save & Deploy** button is grayed out so the change cannot be deployed.

Tunnel Audit

When any type of change is required, whether tunnel modifications or TE resource modifications, a Tunnel Audit is run to determine what inconsistencies the change might cause, if any. Tunnel Audit can also be used anytime to check for the optimality of network utilization.

The audit can be performed from the primary tunnel window or from the TE Links List window. (See [TE Resource Management, page 8-20](#).)

To perform an audit on the created tunnel, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **Create Managed Tunnel**.

The TE Managed Primary Tunnels SR window appears.

Tunnel Audit can be used in two ways:

- When one or more tunnels have been created or their attributes altered (see [Create Primary Tunnel, page 8-46](#)), Tunnel Audit can be activated by selecting **Proceed with Changes**.
- When no changes have taken place, Tunnel Audit can be accessed by selecting **Placement Tools**.

As an example, assume that a new primary tunnel SR has been created.

The TE Managed Primary Tunnel SR window appears.

Step 3 Choose **Proceed with Changes > Tunnel Audit**.

The Computation In Progress window appears temporarily. Then the TE Primary Tunnel Computation Results - Changes window appears.

This window includes the following elements:

Status section (top):

- **Computation Status**—Indicates whether the computation succeeded or failed.
- **Tunnels:**
 - **unplaced**—Number of unplaced tunnels out of the total.
 - **moved**—Number of tunnels that were moved.
- **Bandwidth - unplaced**—Amount of tunnel bandwidth that was not placed out of the total bandwidth of all existing and new tunnels.
- **Global Util.**—Global Pool bandwidth utilization percentage.

The utilization values can be the following:

- **Global Pool**—Comparison data for various Global Pool attributes.
- **Sub Pool**—Comparison data for various Sub Pool attributes.
- **Median**—Utilization of the link that is the middle link when all links are ordered by utilization.
- **Max. Modifiable**—Utilization value for the most utilized link that has movable tunnels passing through it.
- **Mean**—Average link utilization for the network as a whole.
- **Max.**—Utilization value for the most utilized link in the topology.
- **Sub Pool Util.**—Sub Pool bandwidth utilization percentage.
- **Solution**—Utilization for the generated solution.
- **Original**—Utilizations for the original placement.

Changes section (left):

- **Changes**—Number of changes achieved out of the total number of changes.
 - **Achieved**—Indicates whether a specific change is successful (**Yes** or **No**).
 - **Origin**—The originator of the change. Can be **user** (change by user) or **compute** (from a computation, e.g. rerouting of a tunnel).
 - **Type**—The type of change requested: **Tunnel Add Change**, **Tunnel Modify Change**, **Tunnel Remove Change**, or **Element Modify Change**.
 - **Object ID**—A tunnel or link ID.



Note

Certain attributes, such as Description, that do not impact the computation carried out by the placement tools and updates to these are not displayed in the computation results window.

Step 4 To obtain detailed information about the tunnel and whether the change request was achieved, select the specific tunnel and click **Details**.

A **qualityReport** is always generated. If the computation was successful, this will be the only report.

If a warning or a violation was encountered, one or more warning or violation reports will also be generated.

Step 5 To view an audit report, click **View Report**.

In some cases, both a **qualityReport** and a violation report is generated.

Step 6 To view the contents of the **qualityReport**, select the **qualityReport** and click the **Details** button.

The qualityReport fields in the right window pane include the following elements:

Status section (top): described above.

Report section (left):

- **Report Type**—There are three basic report types: a **qualityReport** (generated every time), warning reports, and violation reports.
- **Summary Info**—Summary information about the findings of the report.

Information section (right):

- **Report Type**—See description above.
- **Description**—Specific information about the report.
- **Achievement**—Success or failure of the computation attempt/solution (**SUCCESS** or **CONSTRAINT_VIOLATIONS_REPORTED**).
- **Solution**—Indicates whether a solution was found (**SOLUTION_FOUND**, **PARTIAL_SOLUTION_FOUND** or **NO_SOLUTION_FOUND**).
- **Termination**—Indicates whether the computation was completed:
 - **COMPLETED**—The computation completed processing before the time limit.
 - **TIMED_OUT**—The computation was not able to complete processing within the time limit. The solution presented is the best solution it was able to find in the time available.
- **Optimality**—Indicates whether the computation was optimal:
 - **OPTIMAL_FOR_ALL_CRITERIA**—The solution generated has proven to be the best for all optimization criteria.
 - **NO_OPTIMALITY_PROOF**—The solution's optimality is unknown.
 - **OPTIMAL_FOR_DEMAND_SELECTION**—The solution generated has proven to be the best in terms of total bandwidth placed, but utilization optimality is unknown.

OPTIMAL_FOR_SUB_POOL_PATH_SELECTION—The solution generated has proven to be the best in terms of total bandwidth placed and maximum sub pool utilization, but has not proven to be optimal in terms of global pool utilization.

Step 7 To view the contents of the violation report, select the violation report and click the **Details** button.

The TE Primary Tunnel Computation Results - Report (Details) window appears.

The report fields in the right window pane are described for each report in [Warnings and Violations, page 8-98](#)

Step 8 Click **View Result** to return to the Changes window.

If the proposed changes were achieved, you can click on **Save & Deploy** to save the achievable changes to the repository and implement the tunnel modifications on the network.



Note

Save & Deploy will discard any changes that were not achievable.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Tunnel Placement

The Placement feature supports the admission of new tunnels into the network and the modification of tunnels already admitted into the network. Prime Provisioning will attempt to implement the changes in such a way that network utilization is optimized.

To place a created tunnel, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **Create Managed TE Tunnel**.

The TE Managed Primary Tunnels SR window appears.

Step 3 When one or more tunnels have been created or their attributes altered (see [Create Primary Tunnel, page 8-46](#)), select **Proceed with Changes > Tunnel Placement**.

The Movable Tunnel Selection (Placement) window appears.

Step 4 Set the movable and unmovable managed tunnels.

You can specify whether, when admitting a new tunnel, existing tunnels can be moved (rerouted). This is configurable by you. The default is that managed tunnels are not movable.

Step 5 Click **Proceed**.

The Computation In Progress window shown appears temporarily. Then the TE Primary Tunnel Computation Results - Changes window appears.



Note

Certain attributes, such as Description, that do not impact the computation carried out by the placement tools and updates to these are not displayed in the computation results window.

Step 6 To obtain detailed information about the tunnel and whether the placement request was achieved, select the specific tunnel and click **Detail**.

The detail section in the right side of the window appears.

If the placement request succeeded (**Achieved: yes**), the Detail pane will contain a computed **Path** that is selectable.

To view the path information, click the blue link in the computed **Path** field. The TE Explicit Path window appears.

Step 7 To view the placement report(s), click **View Report** in the Changes window.

The TE Primary Tunnel Computation Results - Report window appears.

A **qualityReport** is always generated. If the computation was successful, this will be the only report.

If a warning or a violation was encountered, one or more warning or violation reports will be generated as well.

Step 8 To view the contents of a placement report, select one of the reports and click the **Details** button.

In the case of a **qualityReport**, the TE Primary Tunnel Computation Results - Report (details) window appears in the report pane on the right.

- Step 9** Click **View Result** to return to the Changes window and click **Save & Deploy** to save the change to the repository and implement the tunnel modifications on network.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Tunnel Repair

As changes are made to the bandwidth requirements or delay parameters of existing tunnels, inconsistencies can arise with the Tunnel Placement. You can run a Tunnel Repair to address such inconsistencies. The objective of Tunnel Repair is to try to move as few existing tunnels as possible to accommodate the changes.

The repair operation can be performed from the primary tunnel window or from the TE Links List window. (See [TE Resource Management](#), page 8-20.)

In the following, we will seek to repair an edited tunnel:

- Step 1** Choose **Traffic Engineering > Create Managed Tunnel**.

The TE Managed Primary Tunnels SR window appears.

Tunnel Repair can be used in two ways:

- When one or more tunnels have been created or their attributes altered (see [Create Primary Tunnel](#), page 8-46), Tunnel Repair can be activated by selecting **Proceed with Changes > Tunnel Repair**.
- When no changes have taken place, Tunnel Repair can be accessed by selecting **Placement Tools > Tunnel Repair**.

- Step 2** In this example, a new primary tunnel SR has been created.

Run Tunnel Repair on the modified tunnels from the TE Managed Primary Tunnels SR window by navigating

Proceed with Changes > Tunnel Repair

The Movable Tunnel Selection window appears.

- Step 3** Set the tunnels that should be movable.

Tunnel Repair will only move existing tunnels if it has to. If you do not want certain tunnels to be moved during Tunnel Repair, these tunnels should be explicitly excluded from the selected list of movable tunnels.

You can also specify a limit on the maximum number of tunnel moves that are acceptable using the **Maximum number of tunnel moves** field.



Note It is not necessary to set modified tunnels to be movable as these are movable by default.

- Step 4** Click **Proceed**.

The Computation In Progress window shown appears temporarily. Then the TE Primary Tunnel Computation Results - Changes window appears.

**Note**

Certain attributes, such as Description, that do not impact the computation carried out by the placement tools and updates to these are not displayed in the computation results window.

Step 5 To obtain detailed information about the tunnel and whether the change request was achieved, select the specific tunnel and click **Detail**.

The detail section in the right side of the window appears.

Step 6 To view a repair report, click **View Report**.

The TE Primary Tunnel Computation Results - Report window appears.

A **qualityReport** is always generated. If the computation was successful, this will be the only report.

If a warning or a violation was encountered, one or more warning or violation reports will also be generated.

Step 7 To view the contents of the repair report, click the **Details** button.

In the case of a **qualityReport**, the TE Primary Tunnel Computation Results - Report (details) window appears.

The report fields in the right window pane are described for each report in [Warnings and Violations, page 8-98](#)

Step 8 Click **View Result** to return to the Changes window and click **Save & Deploy** to save the change to the repository and implement the tunnel modifications on network.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Grooming

The purpose of grooming is to analyze the tunnel pathing with respect to the network elements and optimize resource allocation.

Grooming is not available when change requests have been created. In that case, only the placement tools under **Proceed with Changes** are available.

To perform grooming on the network, use the following steps:

Step 1 Choose **Traffic Engineering > Create Managed TE Tunnel**.

The TE Managed Primary Tunnels SR window appears.

Step 2 Run Grooming by navigating

Placement Tools > Groom

The Movable Tunnel Selection window appears.

Step 3 Set the tunnels that should be movable.

As with Tunnel Repair, Grooming will only move existing tunnels if it has to. If you do not want certain tunnels to be moved during the Grooming process, these tunnels should be explicitly excluded from the selected list of movable tunnels.

Step 4 Click **Proceed**.

The Computation In Progress window shown appears temporarily. Then the TE Primary Tunnel Computation Results - Changes window appears.



Note

Certain attributes, such as Description, that do not impact the computation carried out by the placement tools and updates to these are not displayed in the computation results window.

Step 5 To obtain detailed information about the Grooming and whether it succeeded, select the specific tunnel and click **Detail**.

The detail section in the right side of the window appears.

Step 6 To view a Grooming report, click **View Report**.

The TE Primary Tunnel Computation Results - Report window appears.

A **qualityReport** is always generated. If the computation was successful, this will be the only report.

If a warning or a violation was encountered, one or more warning or violation reports will also be generated.

Step 7 To view the contents of the Grooming report, click the **Details** button.

In the case of a **qualityReport**, the TE Primary Tunnel Computation Results - Report (details) window appears.

The report fields in the right window pane are described for each report in [Warnings and Violations, page 8-98](#)

Step 8 Click **View Result** to return to the Changes window and click **Save & Deploy** to save the change to the repository and implement the tunnel modifications on the network.

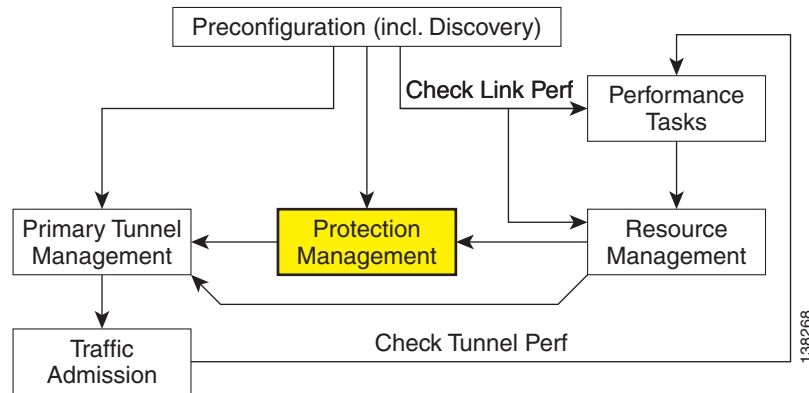
The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Protection Planning

This section describes the process of creating and managing the protection of network elements using automated protection tools. See [Basic Tunnel Management, page 8-27](#) for a description of the process using the basic tools.

The highlighted box in [Figure 8-18](#) shows where in Prime Provisioning protection management occurs.

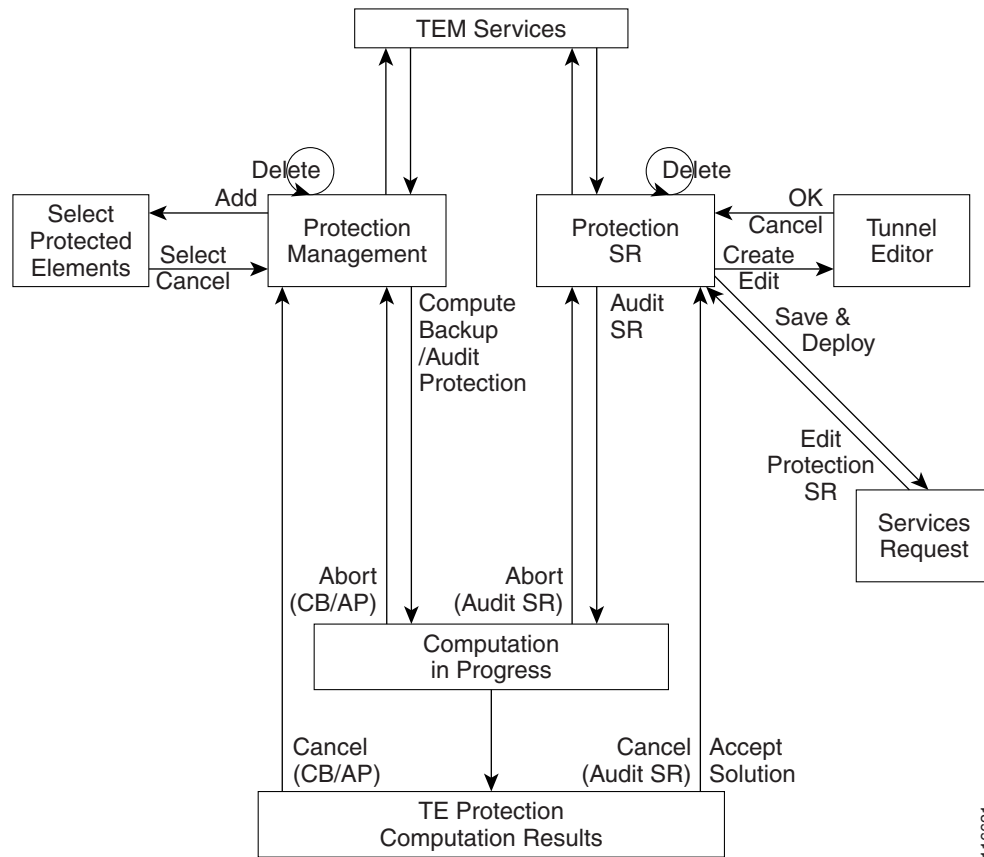
Figure 8-18 Prime Provisioning Process Diagram - Protection Management

The purpose of protection planning is to protect selected elements in the network (links, routers, or SRLGs) against failure.

The first step is to identify the elements that must be protected and then invoke the protection tools to compute the protected tunnels. From the computation, the system responds for each element with either a set of tunnels that protect the element or a set of violations and warnings that help you determine why it could not be protected.

For successfully protected elements the tunnels can be deployed on the network. For elements that could not be protected, the protection is either ignored or the constraints are altered on the protection case. More specifically, this can involve changing the TE bandwidth settings of the links associated to the element and then rerunning the protection computation on the altered network.

An overview of the protection management processes is provided in [Figure 8-19](#).

Figure 8-19 Protection Management Processes

116621

This section includes the following:

- [SRLG Operations, page 8-61](#)
 - [Create SRLG, page 8-61](#)
 - [Edit SRLG, page 8-61](#)
 - [Delete SRLG, page 8-62](#)
- [Configure Element Protection, page 8-62](#)
- [Protection Tools, page 8-62](#)
 - [Compute Backup, page 8-63](#)
 - [Recompute Backup, page 8-64](#)
 - [Audit Protection, page 8-65](#)
 - [Audit SR, page 8-66](#)

SRLG Operations

It is not uncommon for links to have identical physical characteristics, such as being physically located in the same conduit, or being connected to the same hardware. As a result, they could fail as a group during a single failure event. A Shared-Risk Link Group (SRLG) addresses this problem by identifying links that could fail together.

After SRLG modifications (create, edit, delete), use the protection planning functions in the **TE Protection Management** window to ensure that adequate protection is available on the network.

Create SRLG

Creating an SRLG is only necessary if a shared risk link group has been identified and it must be protected.

To create an SRLG, use the following steps:

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Choose Traffic Engineering > SRLGs .
The TE SRLG List window appears. |
| Step 2 | To create an SRLG in the TE SRLG List , click Create .
The TE SRLG Editor window appears. |
| Step 3 | Specify an SRLG Name . |
| Step 4 | Click Add Link .
The Links associated with SRLG window appears. |
| Step 5 | Select one or more links and click Select .
The corresponding link information is added to the link list and the Select window closes and returns to the SRLG editor. |
| Step 6 | Click Save to save the SRLG.
This closes the SRLG editor and brings back the TE SRLG List as the active window, where the newly created SRLG is listed. |
-

Edit SRLG

To edit an SRLG, use the following steps:

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Choose Traffic Engineering > SRLGs .
The TE SRLG List window appears. |
| Step 2 | To edit an SRLG in the TE SRLG List, from the TE SRLG List window select the SRLG that you want to modify and click Edit .
The TE SRLG Editor window appears. |
| Step 3 | Use Add Link and Remove Link to adjust to the desired set of links for the selected SRLG. |
| Step 4 | Click Save to save the changes. |
-

Delete SRLG

To delete an SRLG, use the following steps:

-
- Step 1** Choose **Traffic Engineering > SRLGs**.
The TE SRLG List window appears.
- Step 2** To delete an SRLG in the TE SRLG List, from the TE SRLG List window select the SRLG(s) that you want to delete and click **Delete**. The Delete Confirm window appears.
- Step 3** Click **Delete** to confirm.
The Delete Confirm window closes. After the TE SRLG List window has been updated, the deleted SRLG no longer appears in the SRLG list.
-

Configure Element Protection

Before a protection computation can be performed, it is necessary to configure the network element protection.

To do so, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Protected Elements**.
The TE Protection Management window appears.
Explanation of the **Protection Status** field:
Protection Status—The protection status displayed is determined from the last time an audit was performed. The audit is performed either explicitly by the user or when the protection SR is deployed. The protection status is stated for each network element as either **Protected**, **Not Fully Protected**, or **Unknown**. Click on the column header, **Protected**, to sort elements according to protection status
- Step 2** First, decide which network elements must be protected.
In the TE Protection Management window, click **Add** to add a protection element (link, node, or SRLG). The Select Protection Elements window appears.
Links that are connected to non-Cisco devices cannot be protected and will, therefore, not show in the Select protection elements window. Likewise, non-Cisco devices and SRLGs that contain links to non-Cisco devices cannot be protected and are excluded from the selection.
- Step 3** Select one or more elements to be protected and click **Select**.
The Select Protection Element window closes and the TE Protection Management window reappears.
Next, decide which protection tools should be applied. These are described in [Protection Tools](#), page 8-62.
-

Protection Tools

Relying on manual creation of backup tunnels as described in [Basic Tunnel Management](#), page 8-27 has its limitations, not just for larger and more complicated networks.

The protection tools available in Prime Provisioning provide a number of tools that automatically compute and verify protection of specified network elements.

**Note**

Certain attributes, such as Description, that do not impact the computation carried out by these tools and updates to these are, therefore, not displayed in the computation results window.

Compute Backup

Compute Backup is used to let Prime Provisioning automatically compute the necessary backup tunnels to protect specified network elements. The manual process is described in [Basic Tunnel Management, page 8-27](#)

To run Compute Backup, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Protected Elements**.
 - Step 2** Configure the necessary protection elements as described in [Configure Element Protection, page 8-62](#).
 - Step 3** If you only want to perform Compute Backup on selected elements, select one or more elements on which to calculate a backup path.
 - Step 4** Click **Compute Backup** and select one of the following:
 - All Elements
 - Selected Elements

First the Computation In Progress window appears and then the TE Protection Computation Results window appears.

The **Element:** table displays the outcome of the computation for each element in the protection computation. The status for each element is indicated by at least one row per element in the table. If the status is not valid, the table will contain one row per warning or violation.

The **Element:** table contains the following columns:

- **Element Name**—Name of the network element to be protected.
- **Type**—Network element type (node, link, or SRLG).
- **Report**—Warning or violation associated with an element, if any, as reported by the computation engine.
- **Status**—Computation status of the network element:
 - Valid Tunnels—The element is fully protected by backup tunnels.
 - InvalidTunnels—An Audit Protection detected that the element was not fully protected by the existing backup tunnels.
 - No Solution Exists—A Compute Backup has proven that it is not possible to fully protect the element.

**Note**

Certain attributes, such as Description, that do not impact the computation carried out by the protection tools and updates to these are not displayed in the computation results window.

- Step 5** Select a row corresponding to a specific warning or violation and click **Detail** to display a detailed description in the right pane and backup tunnels associated with the selected item in the bottom pane.
- For a description of warnings and violations, see [Warnings and Violations, page 8-98](#)

Explanation of the **Protection Type** column:

- **Protection Type**—Protection side-effect from activating the tunnel. There are three protection types:
 - **Protection tunnels**—Tunnels that can be activated to provide protection for a specified element.
 - **Side-effect tunnels**—Tunnels that are activated to protect a neighboring element, but which are also activated when a specified element fails.
 - **Activated tunnels**—Tunnels that are activated when a specified element fails, and which might or might not provide protection for the specified element or its neighbors.

The **Backup Tunnel** table displays which new protection tunnels are required and any existing tunnels that should be kept or deleted for each element.

Step 6 If the proposed protection solution is acceptable, click **Accept Solution**.

The TE Protection SR window appears with all tunnel additions and deletions computed by the system.

For an explanation of the various window elements, see [Create Backup Tunnel, page 8-39](#).

Optionally, you can make tunnel changes here and then run **Audit SR** to ensure that you have the desired level of protection before you deploy (see [Audit SR, page 8-66](#)).

Step 7 Click **Save & Deploy** to deploy the new tunnel SR to the network.

When you click **Save & Deploy**, Prime Provisioning locks the TE routers effected, which will block any subsequent SRs which use that TE router until the SRs are finished. It is safe to try and deploy other SRs in the system. If there is any conflict with the SR currently being processed, Prime Provisioning will simply ask you to wait until it is complete. To see the state of deployment, go to the Service Requests window under Inventory and Connection Manager or open the Task Manager under Monitoring.



Note

With the exception of TE Traffic Admission SRs, TE SRs are always deployed immediately from the specific TE SR window, not from the **Service Requests** page in **Inventory and Connection Manager**.

The Service Requests window (**Operate > Service Request Manager**) opens and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

If the SR does not go to the **Deployed** state, go to the Task Logs window to see the deployment log (**Monitoring > Task Manager > Logs**) as described in [SR Deployment Logs, page 11-47](#).

Recompute Backup

Recompute Backup is used to automatically recompute existing backup tunnels to update protection for specified network elements that are in either **Protected**, **Not Fully Protected**, or **Unknown** state.

The function Compute Backup Tunnels attempts to minimize changes to existing tunnels. Thus, if it can create new backup tunnels to protect the required elements without making any changes to existing tunnels, it will do so. Minimizing change is useful but also has a disadvantage if new resources such as more links or more bandwidth on the links have been added to the network. The current tunnels would not be changed to take advantage of those new resources even though they might provide better and shorter protection paths.

The Recompute Backup Tunnels function is for just these cases. It will compute backup paths for new and existing FRR tunnels without attempting to maintain the current paths.

To run Recompute Backup, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Protected Elements**.
- Step 2** Select one or more elements on which to recompute a backup path.
- Step 3** Click **Recompute Backup** and select one of the following:
- All Elements
 - Selected Elements

First the Computation In Progress window appears and then the TE Protection Computation Results window.

For a description of these fields, see [Step 4](#) under [Compute Backup](#), page 8-63.

- Step 4** For the rest of the procedure, see Step 5 and onwards in the procedure documented in [Compute Backup](#), page 8-63.
-

Audit Protection

As opposed to the Compute Backup tool described on page 63, Audit Protection does not attempt to create a backup solution. It seeks to verify protection of specified network elements with the current set of backup tunnels and reports any warnings or violations that are discovered. It is recommended that any time a change has been committed on the TE topology such as resources on TE links or SRLG membership, a protection audit be run to verify the protection status on all elements.

The computation will display the same computation results page as for Compute Backup. When you return from the computation results page, the Protection Status column in the TE Protection Management window is updated to show the level of protection for each element.

This section describes the necessary steps to perform Audit Protection on one or more network elements.

To run Audit Protection, use the following steps:

-
- Step 1** Choose **Traffic Engineering > TE Protected Elements**.
- The TE Protection Management window appears.
- Explanation of the **Protection Status** field:
- Protection Status**—The protection status displayed is determined from the last time an audit was performed. The audit is performed either explicitly by the user or when the protection SR is deployed. The protection status is stated for each network element as either **Protected**, **Not Fully Protected**, or **Unknown**. Click on the column header, **Protected**, to sort elements according to protection status
- Step 2** If you only want to perform Audit Protection on selected elements, select one or more tunnels on which to calculate a backup path.
- Click **Audit Protection** and select one of the following:
- All Elements
 - Selected Elements

The Computation In Progress window appears.

Then the TE Protection Computation Results window appears.

For an explanation of the various window elements, see [Compute Backup, page 8-63](#).

**Note**

Certain attributes, such as Description, that do not impact the computation carried out by the protection tools and updates to these are not displayed in the computation results window.

Step 3 To view the backup tunnels for a particular element, select the element and click **Details**.

The TE Protection Computation Results window appears.

For an explanation of the various window elements, see [Compute Backup, page 8-63](#).

Step 4 Select a row corresponding to a specific warning or violation and click **Details** to display a detailed description in the right pane and backup tunnels associated with the selected item in the bottom pane. Tunnels associated with a warning or violation are flagged in the **Report** column in the **Backup Tunnels** table in the bottom pane.

The **Accept Solution** button is greyed out because the audit does not provide a solution but rather an evaluation.

For a description of warnings and violations, see [Warnings and Violations, page 8-98](#)

Step 5 Click **Cancel** to return to the TE Protection Management window.

The protection status is updated in the Protection Status column.

Audit SR

Audit SR audits protection of all elements in the **TE Protection Management** window against backup tunnels in the TE Protection SR window.

This feature can be used to audit the protection for manually added, modified, and deleted tunnels in the TE Protection SR window before deploying them.

To audit a TE backup tunnel SR, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **Create TE Backup Tunnel**.

The **TE Protection SR** window appears. For an explanation of the various window elements, see [Create Backup Tunnel, page 8-39](#).

Step 3 To audit the protection SR, click **Audit SR**.

**Note**

Audit SR will only be enabled if there are elements in the TE Protection Management window. If this is not the case, the **Audit SR** button will be disabled (grayed out).

The FRR Audit process begins and the TE Protection Computation Results window appears.

See [Audit Protection, page 8-65](#) for a description of the rest of the process. Detail and report windows are identical in these two processes.

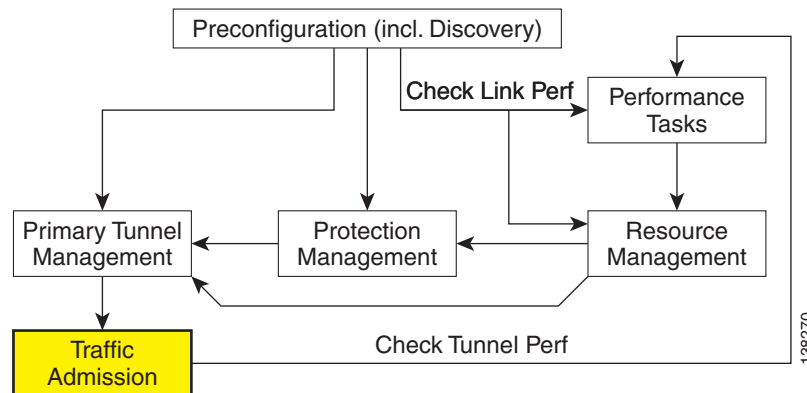
TE Traffic Admission

TE Traffic Admission is the first step towards enabling services on TE tunnels. There are a number of mechanisms that can be used for forwarding traffic into a tunnel to provide basic IP connectivity. The current implementation of Cisco Prime Provisioning Traffic Engineering Management (Prime Provisioning) uses both static routing and autoroute announce to inform the routing protocol of the tunnel's presence. Autoroute announce can be also used as part of the routing protocol calculations.

The TE Traffic Admission tool is used to assign traffic to traffic-engineered tunnels.

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning TE Traffic Admission occurs.

Figure 8-20 Prime Provisioning Process Diagram - TE Traffic Admission



Static routing is perhaps the simplest way of forwarding traffic into a tunnel. Traffic that matches a target destination prefix is routed into a particular tunnel.

While this achieves the basic goal of directing traffic into a given tunnel, this approach has limitations. First, the offering of differentiated Class-of-Service (CoS) treatment is limited to destination-based CoS. As each source PE serves as an aggregation point for a number of traffic flows, there is no way to restrict which traffic receives preferential treatment to a destination because access to a tunnel is through general routing. Secondly, it does not generally provide a scalable solution because the static routing mechanism must capture both the large number of subnets that can be served by each PE router, and it must be able to further capture CoS treatment for each of these subnets.

Static routing works best if there is no need to provide differentiated CoS treatment by destination. That is, all packets destined for one or more particular prefixes all receive the same CoS.

This section includes the following:

- [Creating a TE Traffic Admission SR, page 8-68](#)
- [Deploying a TE Traffic Admission SR, page 8-69](#)
- [Other Traffic Admission SR Operations, page 8-70](#)
- [Viewing the SR State, page 8-70.](#)

Creating a TE Traffic Admission SR

The TE traffic admission tool in Cisco ISC TEM only displays primary tunnels (managed or unmanaged) when they are associated with a TE provider and the tunnels are not already associated with a TE Admission SR. That is, the tool is only intended for admitting new traffic onto tunnels currently not carrying any traffic.

To create a TE Traffic Admission SR, use the following steps:

Step 1 Choose **Traffic Engineering**.

Step 2 Click **TE Traffic Admission**.

The TE Traffic Admission Tunnel Selection window appears.



Note If this window does not open, either no tunnels are associated with a TE provider or any tunnels associated with a TE provider are already tied to a TE Admission SR.

The TE Traffic Admission Tunnel Selection window lists all primary tunnels, both managed and unmanaged, that are not already associated with an admission SR.

The **Deploy Status** can be **Pending**, **Deployed**, or **Functional**.



Note

Backup tunnels are not displayed in the TE Traffic Admission Tunnel Selection window.

Step 3 Select a TE tunnel by clicking the corresponding radio button and clicking **Select**.

The TE Traffic Admission SR window appears.

The main TE Traffic Admission SR window includes the following fields:

- **Tunnel**—Tunnel name.
- **Description**—Service request description.
- **EXP** [IOS devices only]—Class marking bits for CBTS.
- **Policy** [IOS XR devices only]—Policy marking bits for PBTS.
- **Autoroute announce**—Used to specify that the Interior Gateway Protocol (IGP) should use the tunnel (if the tunnel is up) in its enhanced shortest path first (SPF) calculation.
 - **On**—Autoroute announce is enabled.
 - **Off**—Autoroute announce is disabled.
- **Autoroute Metric**—Used to specify the Multiprotocol Label Switching (MPLS) traffic engineering tunnel metric that the Interior Gateway Protocol (IGP) enhanced shortest path first (SPF) calculation uses.
 - **Absolute**—Absolute metric mode; you can enter a positive metric value.
 - **Relative**—Relative metric mode; you can enter a positive, negative, or zero value.
- **Static Routes**—Lists any static routes that the tunnel uses.
- **Destination**—Name of the static route for the tunnel destination.
- **Distance**—Administrative distance (cost).

**Note**

If TE Traffic Admission SR attributes such as PBTS attributes are changed outside Prime Provisioning and a TE discovery task is run, the discovery task logs will not report a discrepancy warning and the repository will be updated with the new configuration from the device.

Step 4 When filling out the form, if **Autoroute Announce** is set to **On**, indicate whether **Autoroute Metric** should be **Absolute** or **Relative**.

Step 5 You can also set an optional autoroute metric.

For the relative metric, the range is -10 to 10, for the absolute metric, the range is 1 to 2147483647.

**Note**

CBTS is supported in IOS and PBTS is supported in IOS XR. If the tunnel head router is running IOS XR, the **EXP** fields will not be present and are replaced with the **PBTS** fields.

When clicking the **Add** button, the Add TE Static Route window appears.

Step 6 In the Add TE Static Route window, specify at a minimum a **Destination IP** address (w.x.y.z/n). Optionally specify an administrative **Distance**. It is recommended that you either define one or more static routes or, alternatively, that you define an autoroute.

Step 7 Click **OK** to accept the entries or **Cancel** to exit the window.

In the main TE Traffic Admission SR window, you can add another TE Static Route or edit existing routes.

Step 8 Click **Save** to save the service request.

The Service Requests window appears with the TE Traffic Admission SR in **REQUESTED** state and the Operation Type set to **ADD**.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

To deploy the service request from the Service Requests window, see [Deploying a TE Traffic Admission SR, page 8-69](#).

Deploying a TE Traffic Admission SR

As opposed to the TE Primary Tunnel SR, Backup Tunnel SR, and TE Resource Modification windows, a TE Admission SR must be deployed from the general Service Request Manager window.

To deploy a TE Admission SR, use the following steps:

Step 1 Choose **Operate > Service Request Manager**.

The Service Requests window appears.

The Service Requests window includes the following elements:

- **Job ID**—Job ID for the SR.
- **Data Files**—This field is used for variable substitutions via templates and currently do not apply to TEM SRs.

- **State**—Indicates whether the tunnel state is **DEPLOYED** or **NOT DEPLOYED** and whether it is **Conformed** or **Not Conformed**.
- **Type**—The type of service request, indicating which service issued the request. For a detailed description of the possible service types, see the managing service requests part elsewhere in this guide.
- **Operation Type**—SR operation on the tunnel, can be either **ADD**, **MODIFY**, **DELETE**, or **ADMIT**. Applicable only to tunnels in the current SR.
- **Creator**—ID for the user who created the SR.
- **Customer Name**—Name of the customer to which the SR applies.
- **Policy Name**—Name of the policy associated with the SR.
- **Last Modified**—Date and time when the SR was last modified.
- **Description**—SR description provided by the user.

Step 2 Select the desired service request and click **Deploy**.

A drop-down menu appears under the **Deploy** button. In the drop-down menu, select **Deploy** or **Force Deploy**. After having been successfully deployed, the **State** of the SR changes to **Deployed**.

The Service Requests window (**Operate > Service Request Manager**) appears and displays the state of the deployed SR.

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

Other Traffic Admission SR Operations

As opposed to other service requests, TE Traffic Admission SRs can be decommissioned in the Service Requests window.

Edit and decommission operations for TE Traffic Admission service requests are handled in the Service Request Manager window. These operations are described in the managing service requests part elsewhere in this guide.

Viewing the SR State

To view a service request state, go to **Operate > Service Request Manager**.

If the SR does not enter the **Deployed** state, go to the **Task Logs** window to see the deployment log (**Operate > Task Manager > Logs**) as described in [SR Deployment Logs, page 11-47](#).

Administration

A number of administrative features in Cisco Prime Provisioning Traffic Engineering Management (TEM) are common to Prime Provisioning. Instructions on how to use these features are described in detail starting in [Cisco Prime Provisioning 6.4 Administration Guide](#).

In this section, only TE-specific administrative features are described.

This section includes the following:

- [TE User Roles, page 8-71](#)
- [TE Policies, page 8-71](#)
 - [Create Policy, page 8-71](#)
 - [Edit Policy, page 8-72](#)
 - [Delete Policy, page 8-73](#)
- [TE Tasks, page 8-73](#)
 - [Creating a TE Task, page 8-73](#)
 - [Creating a TE Functional Audit Task, page 8-74](#)
 - [Creating a TE Interface Performance Task, page 8-75](#)
- [SR History and Configlets, page 8-77](#)
- [Managing the Locking Mechanism, page 8-77.](#)

TE User Roles

A TE user role can be a predefined or a user-specified role defining a set of permissions. For a detailed description of user roles in Prime Provisioning and how to use them, see [Cisco Prime Provisioning 6.4 Administration Guide](#).

To access the User Roles window and locate the TE user roles, choose **Administration > Roles**. The User Roles window appears.

There are two pre-defined TEM user roles:

- **TERole**—Grants full permission to TEM operations.
- **TEServiceOpRole**—Grants permission only to manage the TE Admission SR.

TE Policies

Policies are used to define common tunnel attributes. Attributes such as bandwidth pools, hold and setup priority, and affinity bits, are set manually during policy creation as described below.

This section describes the following policy operations:

- [Create Policy, page 8-71](#)
- [Edit Policy, page 8-72](#)
- [Delete Policy, page 8-73](#)

Create Policy

Prime Provisioning allows you to create TE-specific policies in a manner similar to other policies.

To create a TE policy, use the following steps:

-
- | | |
|---------------|----------------------------------------------------------------------------------------------|
| Step 1 | Choose Service Design > Policy Manager .

The Policy Manager window appears. |
| Step 2 | Click Create and select TE from the drop-down list to set up a new TE policy. |

The TE Policy Editor window appears.

It includes the following fields:

- **Policy Name**—Name of the TE policy chosen by the user.
- **Policy Owner**—The owner of the TE policy.
- **Managed**—Check this box to make the policy to be used by managed tunnels. When clicked, both the setup and hold priorities are set to zero and these are not editable. If the box is unchecked, the setup/hold priorities can be set to a value between 1 and 7.

Clicking the **Managed** check box will add some extra fields in the TE Policy Editor corresponding to two additional protection levels for **FRR Protection Level** (Fast Re-Route) and a new field, **Delay Constraint**.

- **Pool Type**—Tunnel bandwidth pool type for this policy. For a definition of pool types, see the Bandwidth Pools section in [Traffic Engineering Management Concepts, page 8-111](#).
 - **Sub Pool (BC1)**—Bandwidth will be reserved from Sub Pool.
 - **Global Pool (BC0)**—Bandwidth will be reserved from Global Pool.
 - **Setup Priority**—Priority used when signaling an LSP for the tunnel to determine, which of the existing tunnels can be preempted. Valid values are from 0 to 7, where a lower number indicates a higher priority. Therefore, an LSP with a setup priority of 0 can preempt any LSP with a non-0 hold priority.
 - **Hold Priority**—Priority associated with an LSP for the tunnel to determine if it should be preempted by other LSPs that are being signaled. Valid values are from 0 to 7, where a lower number indicates a higher priority.
 - **Affinity**—Attribute values required for links carrying the tunnel (bit values are either 0 or 1).
 - **Affinity Mask**—Which attribute values should be checked. If a bit in the mask is 0, a link's attribute value of that bit is irrelevant. If a bit in the mask is 1, the link's attribute value and the tunnel's required affinity for that bit must match.
 - **FRR Protection Level**—Level of Fast Reroute protection required on the primary tunnel.
 - **None**—No backup tunnel needed.
 - **Best Effort**—Use backup tunnel if available.
 - **Link & SRLG**—Primary tunnel must pass through only links or SRLGs that are FRR-protected
 - **Link, SRLG & Node**—Primary tunnel must pass through only intermediate nodes and links or SRLGs that are FRR-protected.
 - **Delayed Constraint**—Apply a constraint when optimizing paths or placing tunnels.
 - **Max. Delay (msec)**—Sets the maximum delay allowed for each managed tunnel in a given policy.
 - **MPLS IP Enabled**—This configures the tunnel with the **mpls ip** command if enabled.
-

Edit Policy

A policy can be edited only if it is not associated with a tunnel.

To edit a TE policy, use the following steps:

-
- Step 1** Choose **Service Design > Policy Manager**.
The Policies window appears.
- Step 2** Select the desired policy and click **Edit**.
The TE Policy Editor window appears. The policy editor is described in [Create Policy, page 8-71](#). The only difference between the create and edit processes is that the policy name and owner are not editable when editing a policy.
- Step 3** Make the desired changes to the policy attributes and click **Save**.
If the save operation succeeds, the new TE policy now appears in the Policies window. If not, the **Status** box will indicate the type of error that occurred and, when possible, the corrective action required.
-

Delete Policy

A policy can be deleted only if it is not associated with a tunnel.

To delete a TE policy, use the following steps:

-
- Step 1** Choose **Service Design > Policy Manager**.
The Policies window appears.
- Step 2** Select the desired policy and click **Delete**.
The Confirm Delete window appears.
- Step 3** Check the policy marked for deletion and click **OK**.
The Policies window refreshes and the selected policy disappears.
-

TE Tasks

Prime Provisioning currently offers three TE-specific tasks that are used in a manner similar to other tasks:

- **TE Discovery (Full and Incremental)**—Populates the repository with data from the TE network. Discrepancies are reconciled and/or reported.
- **TE Functional Audit**—Performs functional audit on TE Primary or Backup SRs in certain states.
- **TE Interface Performance**—Calculates the interface/tunnel bandwidth utilization.

This section focuses on describing how to create TE Functional Audit and TE Interface Performance tasks. Instructions on how to create a TE Discovery task are included in [TE Network Discovery, page 8-10](#).

Creating a TE Task

TE tasks are managed in the **Task Manager**, which is accessed by selecting **Operate > Task Manager**.

The Tasks window appears.

For a detailed description of the window elements in the Tasks window, see [Task Manager, page 11-24](#).

This page shows all collection and deployment tasks that have been executed. Note that a task could be scheduled to happen once or there could be several scheduled runs of a task. The schedule can be viewed by selecting a task and clicking **Schedules**.

Creating a TE Functional Audit Task

For each tunnel in the SR, the TE Functional Audit task checks the LSP currently used on a router against the LSP stored in the repository:

- tunnel down—Ignore (do not check)
- tunnel up—Check the LSP used on the router against the one stored in the repository:
 - If they are the same, the tunnel and the SR are both set to **Functional**.
 - If they are different, both the tunnel and the SR are set to **Broken**.
- tunnel missing from router—SR left untouched. The tunnel state is set to **Lost**.

This task only performs functional audit on TE Primary or Backup SRs, which are not in one of the following states:

- **Closed**
- **Requested**
- **Invalid**
- **Failed Deploy**

For more information on working with service requests, see the managing service requests part elsewhere in this guide.

To create a TE Functional Audit task, use the following steps:

-
- Step 1** Choose **Operate > Task Manager**.
- Step 2** Click **Audit > TE Functional Audit** to open the Create Task window.
For a detailed description of the window elements in the Create Task window, see [Task Manager, page 11-24](#).
- Step 3** Modify the **Name** or **Description** fields as desired and click **Next**.
The Task Service Requests window appears.
- Step 4** Click **Add** to add a task service request.
The Select Service Request(s) window appears.
- Step 5** Select an SR using the **Select** button.



Note Only SRs of type TE Tunnel or TE Protection will be accepted.

The Selected Service Request(s) window closes and the selected task(s) now appears in the Task Service Requests window. To add other SRs, repeat the procedure in [Step 4](#) and [Step 5](#).

- Step 6** In the Task Service Requests window, click **Next**.
The Task Schedules window appears.
- Step 7** Click **Now** to start the task immediately or **Create** to create a task schedule.
When selecting **Now**, a line is added to the **Task Schedules** window. When selecting **Create**, the Task Schedule window appears.

Step 8 In the Task Schedule window, indicate when and how often to run the task.

Step 9 Click **OK**.

The scheduled task should now appear in the **Task Schedules** table.



Note

The default setting is to schedule a single TE Functional Audit task to take place immediately (“**Now**”).

Step 10 Click **Next**.

The Task Schedule window now shows the new task in its list of created tasks. A summary of the scheduled task appears.

Step 11 Click **Finish**.

This adds the task to the list of created tasks in the Tasks window.

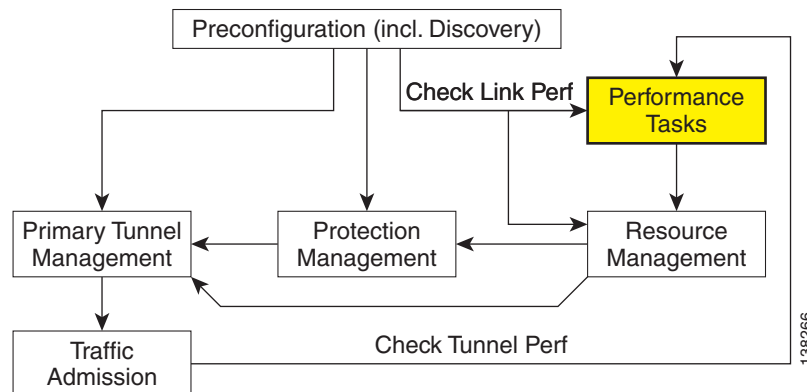
To view the task logs for the created tasks, see [Viewing a Task Log, page 11-47](#).

Creating a TE Interface Performance Task

This task calculates interface/tunnel bandwidth utilization using the Simple Network Management Protocol (SNMP).

The highlighted box in [Figure 8-3](#) shows where in Prime Provisioning traffic admission occurs.

Figure 8-21 Prime Provisioning Process Diagram - TE Interface Performance



Calculating utilization depends on how data is presented for the object you want to measure. Interface utilization is the primary measure used for network utilization. Because MIB-II variables are stored as counters, you must take two poll cycles and figure the difference between the two (hence, the delta used in the equation).

Three variables are required:

- task duration—how long the task will run (in seconds)
- frequency—how frequent the data will be collected (in seconds)
- interval—the distance between two poll cycles (in milliseconds).

The following explains the variables used in the formulas:

- delta(traffic in)—the delta between two poll cycles of collecting the SNMP input object, which represents the number of inbound units of traffic
- delta(traffic out)—the delta between two poll cycles of collecting the SNMP output object, which represents the number of outbound units of traffic
- bandwidth—the speed of the interface.

A more accurate method is to measure the input utilization and output utilization separately, using the following formula:

$$\text{delta(traffic in)} \times 8 \times 100$$

$$\text{Input utilization} = \frac{\text{delta(traffic in)} \times 8 \times 100}{(\text{number of seconds in delta}) \times \text{bandwidth}}$$

$$\text{delta(traffic out)} \times 8 \times 100$$

$$\text{Output utilization} = \frac{\text{delta(traffic out)} \times 8 \times 100}{(\text{number of seconds in delta}) \times \text{bandwidth}}$$

To create a TE Interface Performance task, use the following steps:

-
- Step 1** Choose **Operate > Task Manager**.
- Step 2** Click **Create > TE Interface Performance** to open the Create Task window for a new TE Interface Performance task.
- For a detailed description of the window elements in the Create Task window, see [Task Manager, page 11-24](#).
- Step 3** Modify name and description if needed and click **Next**.
- The Select TE Provider window appears.
- Step 4** Click a radio button to select a TE provider.
- Step 5** Click **Next**.
- The TE Performance Collection window appears.
- Step 6** Enter desired values in the **Task Duration**, **Task Frequency**, and **Task Interval** fields.



Note

If the **Task Interval** field is set too low, the MIB might not be updated, in which case the TE Performance Report will not show any traffic. For tunnels or links on IOS routers, it is recommended to set the interval to 1000 ms; for IOS XR routers, a recommended interval is 5000 ms. Note that these values might need to be tuned to suit your specific environment.

-
- Step 7** Use the **Add** button to select a tunnel or link on which to run the interface performance task:
- **TE Tunnel**—Add a TE tunnel. The Select Tunnel(s) window appears.
 - **TE Link**—Add a TE link. The Select Link(s) window appears.
- Step 8** Select one or more of tunnels and links and click **Next**.
- The selected tunnels and links are added to the **Targets** list in the TE Performance Collection window. The Task Schedules window appears.
- Step 9** Click **Now** or **Create** to create a task schedule.

When you select **Create** to customize the schedule, the Task Schedule window appears (with **Now**, this step is skipped).



Note The default setting is to schedule a single TE Interface Performance task to take place immediately (“**Now**”).

Step 10 In the Task Schedule window, make your selections to define when and how often to run the task.

Step 11 Click **OK**.

The scheduled task should now appear in the **Task Schedules** table.

Step 12 Click **Next**.

A summary of the scheduled task appears.

Step 13 Click **Finish**.

This adds the task to the list of created tasks in the Tasks window.

To view the TE Performance Report that is generated for TE Interface Performance task(s), see [TE Performance Reports, page 11-48](#).

To view the task logs for the created tasks, see [Viewing a Task Log, page 11-47](#).

SR History and Configlets

The history and configlets associated with individual service requests can be viewed from the Service Requests window when you select a service request and click the **Details** button.

The history of a service request is essentially a state change report. It lists the various states that elements associated with an SR has transitioned between and reports relevant details pertaining to these state changes.

Configlets for devices associated with service requests are in simple scrollable text format.

For more information about these features and how to manage service requests, see the managing service requests part elsewhere in this guide.

Managing the Locking Mechanism

Whenever a task is performed that incurs a database update, which might affect the resource and hence the result of a tunnel computation, it locks the system before the update and releases it at completion of the update. If for some reason the lock is not released, other updates that require the lock are blocked.

The purpose of the lock feature is to prevent concurrent and mutually inconsistent planning activities from being committed to the database. Meaning, if each user takes the same snapshot of the the repository, performs computations, and tries to commit what he/she sees, the locking mechanism helps synchronize the commit and ensures that no commit invalidates other commits.

If the system is locked for prolonged periods of time, the administrator should check if anyone is performing long planning tasks and take note of which process locked the system and report it. If the administrator is sure that no one is using the system, it can be unlocked by using the lock manager.

Prime Provisioning has two kinds of locks:

- TE provider lock—Locks managed tunnels, backup tunnels, resource SRs, and TE Discovery.
- TE router lock—Locks unmanaged tunnels.

Each system lock is linked to a TE provider. In the following, procedures for unlocking each system lock are listed.

Unlocking the TE Provider Lock

To unlock the TE provider, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Providers**.
The TE Providers window appears.
- Step 2** Select a TE provider that is locked by checking the corresponding check box.
- Step 3** Click **Manage Lock**.
The System Lock Management window appears.
The text fields in this window are read-only.
- Step 4** To unlock, click the **Unlock** button.
The System Lock Management window closes and the **System Lock Status** field in the TE Providers window is updated accordingly.
-

Unlocking the TE Router Lock

To unlock the TE router lock, use the following steps:

-
- Step 1** Choose **Traffic Engineering > Nodes**.
The TE Nodes List window appears.
- Step 2** Select a TE node that is locked by clicking the corresponding check box.
- Step 3** Click **Manage Lock**.
The System Lock Management window appears. The text fields in this window are read-only.
- Step 4** To unlock, click the **Unlock** button.
The System Lock Management window closes and the **System Lock Status** field in the TE Nodes List window is updated accordingly.
-

Locking Operation Errors

TEM locks the TE Provider or TE Router object respectively for the duration of a save and deploy operation to ensure database consistency.

This section describes the following errors:

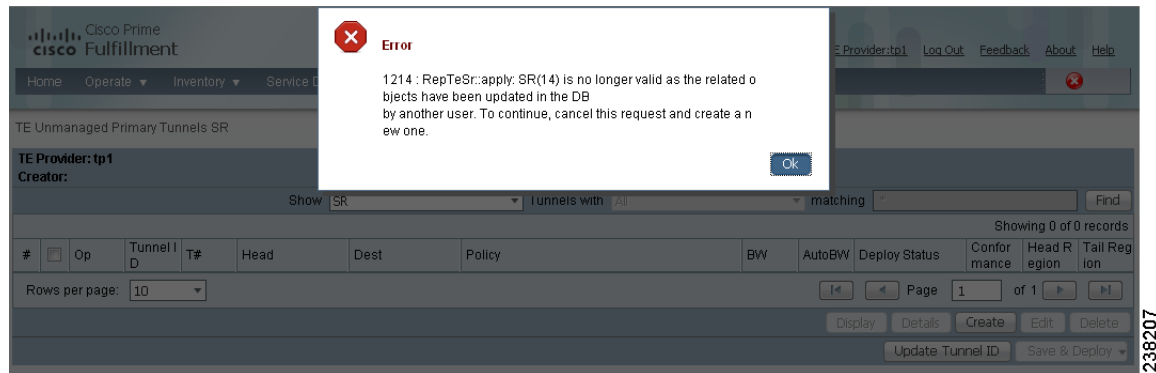
- [Modifying Locked Object, page 8-79](#)
- [Modifying Object After Lock Is Released, page 8-79](#)

- [Deleting Link with Associated TE Object, page 8-79](#)
- [Deleting Link Without Associated TE Object, page 8-80](#)

Modifying Locked Object

If you attempt to modify a locked object, you will be informed that the object cannot be modified because another user is making changes. You will receive the error message shown in [Figure 8-22](#).

Figure 8-22 **Modifying Locked Object**

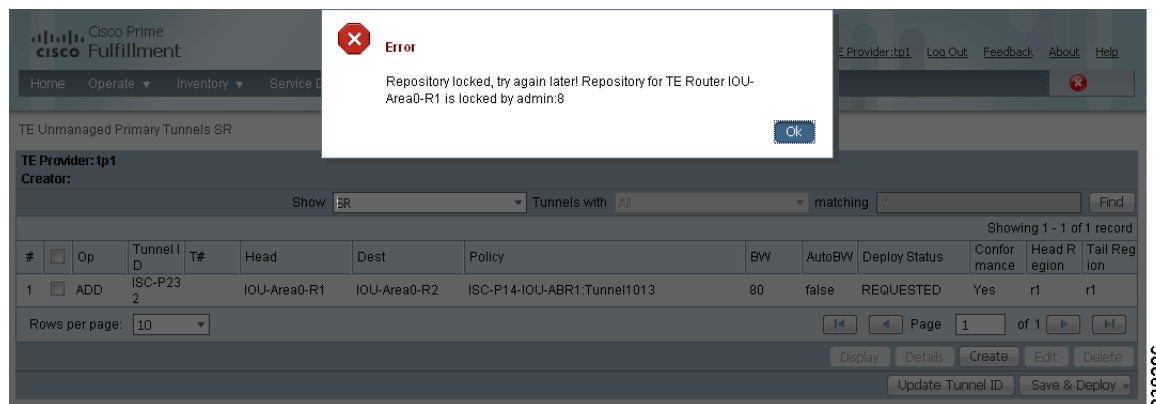


238207

Modifying Object After Lock Is Released

If you attempt to modify an object after the lock is released, Prime Provisioning will check that your current working version of the object is up to date. If not, you will be instructed to restart with a new version of the object as your data is now out of date. You will receive the error message shown in [Figure 8-23](#).

Figure 8-23 **Modifying Object After Lock Is Released**



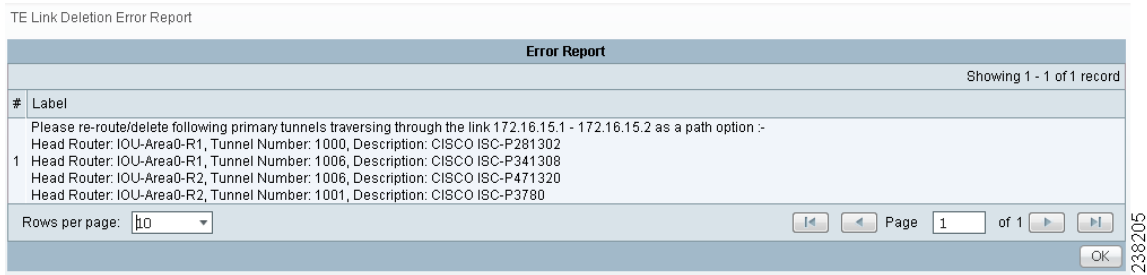
238206

Deleting Link with Associated TE Object

Link removal is not allowed if the link is associated with an explicit path or is traversed by a tunnel.

If you try to delete a link with one or more associated objects, the error message in [Figure 8-24](#) is displayed.

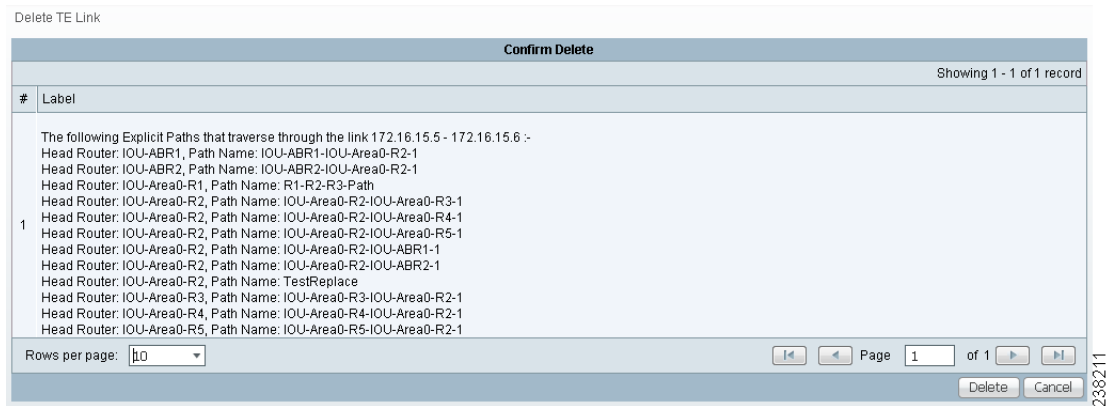
Figure 8-24 Deleting Link with Associated TE Object



Deleting Link Without Associated TE Object

A link can be removed if it is not traversed by a tunnel, even if it is associated with an explicit path. When you try to delete such a link, the type of report shown in [Figure 8-25](#) will be displayed.

Figure 8-25 Deleting Link Without Associated TE Object



TE Topology

The TE Topology tool provides a graphical view of the network set up through the Cisco Prime Provisioning web client. It gives a graphical representation of the various network elements, including devices, links, and tunnels. It also displays devices that Prime Provisioning is unable to identify but which have been discovered with the TE Discovery tool to be part of the network.

The TE Topology tool is accessed from the Traffic Engineering menu.

The TE Topology tool is used to visualize the TE network based on the data contained in the repository. To that end, it provides a number of ways of manipulating the display, for example by applying algorithms to the graph layout, importing maps, and so on.

The tool is accessed from a TE Topology Interface Applet that displays the TE topology through a Java applet within the browser.

This section describes how to use the topology tool.

It includes the following sections:

- [Using the TE Topology Interface Applet, page 8-81](#)
 - [Displaying and Saving Layouts, page 8-83](#)
 - [Using Maps, page 8-84](#)
 - [Using Highlighting and Attributes, page 8-86](#)
 - [Using Algorithms, page 8-87.](#)

Using the TE Topology Interface Applet

The TE Topology Interface Applet (Topology Applet) provides a means of visualizing the network and tunnels present in the network. The web-based GUI is the primary means of visualizing the network information. The Topology Applet simply augments the web-based GUI to provide you with a different presentation format.

The features offered through the Topology Applet are:

- TE Topology rendering
- Highlighting of network elements
- Tunnel overlay (unmanaged, primary, and backup)
- Topology layout persistence
- Integration with web page content.

To access the Topology Applet, use the following steps:

Step 1 Choose **Traffic Engineering > Topology**.

Step 2 Click **TEM Topology Interface Applet**.

If the security certificate for the topology applet has not been accepted previously, you might get a security warning window.

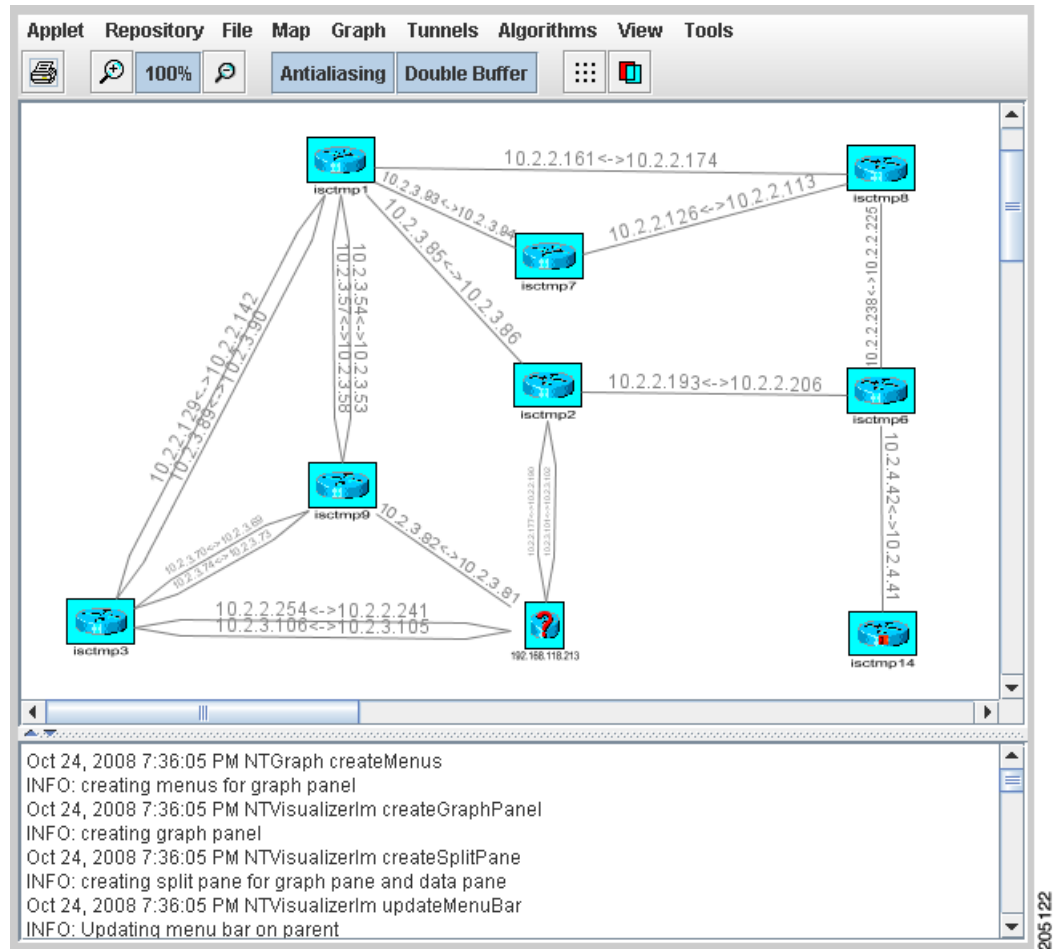
Step 3 Click **Yes** or **Always** to accept the authenticity of the security certificate.

The Topology Display applet window in [Figure 8-26](#) appears.

The screenshot displays the NTVisualizer application window. The menu bar includes: Applet, Repository, File, Map, Graph, Tunnels, Algorithms, View, Tools. The toolbar contains icons for file operations, zooming, and display settings like 'Antialiasing' and 'Double Buffer'. The main area shows a network graph with nodes labeled 'isctmp1' through 'isctmp14'. Edges between nodes are labeled with IP addresses and ranges, such as '10.2.2.238 <-> 10.2.2.225' and '10.2.3.85 <-> 10.2.3.86'. A status bar at the bottom contains a log of events:

- Oct 10, 2008 12:34:39 AM NTGraph createMenus
- INFO: creating menus for graph panel
- Oct 10, 2008 12:34:39 AM NTVisualizerIrm createGraphPanel
- INFO: creating graph panel
- Oct 10, 2008 12:34:40 AM NTVisualizerIrm createSplitPane
- INFO: creating split pane for graph pane and data pane
- Oct 10, 2008 12:34:40 AM NTVisualizerIrm updateMenuBar

After the nodes have been arranged to your liking, you might end up with a topology display similar to the one in [Figure 8-27](#).

Figure 8-27 Topology Display Applet with User-Arranged Topology

Displaying and Saving Layouts

Use the two operations in the **Repository** menu, **Layout Graph** and **Save Graph Layout**, to display or save the current layout of the network graph.

Prior to generating the graph layout, the coordinates must be set on each of the network devices. Otherwise, the graph will have a random layout.

- **Layout Graph**—The graph is laid out from the repository. If a graph layout is already present, that layout is cleared once you click **Yes** in the **Clear Graph Layout** confirmation box. If the layout has not previously been saved, a random layout of the repository contents is drawn. If it has been saved previously, the saved layout is redrawn.
- **Save Graph Layout**—Save the current graph layout. Doing so will ensure that whenever the graph layout is cleared with **Layout Graph** or the topology applet is closed, the same layout will be created when the applet is restarted. If a map was used, the map is also redrawn.

Using Maps

You can associate a map with each view. Currently, the topology viewer only supports maps in the Environmental Systems Research Institute, Inc. (ESRI) shape format. The following sections describe how to load maps and selectively view map layers and data associated with each map.

The map features are accessed from the **Map** menu in the Topology window.

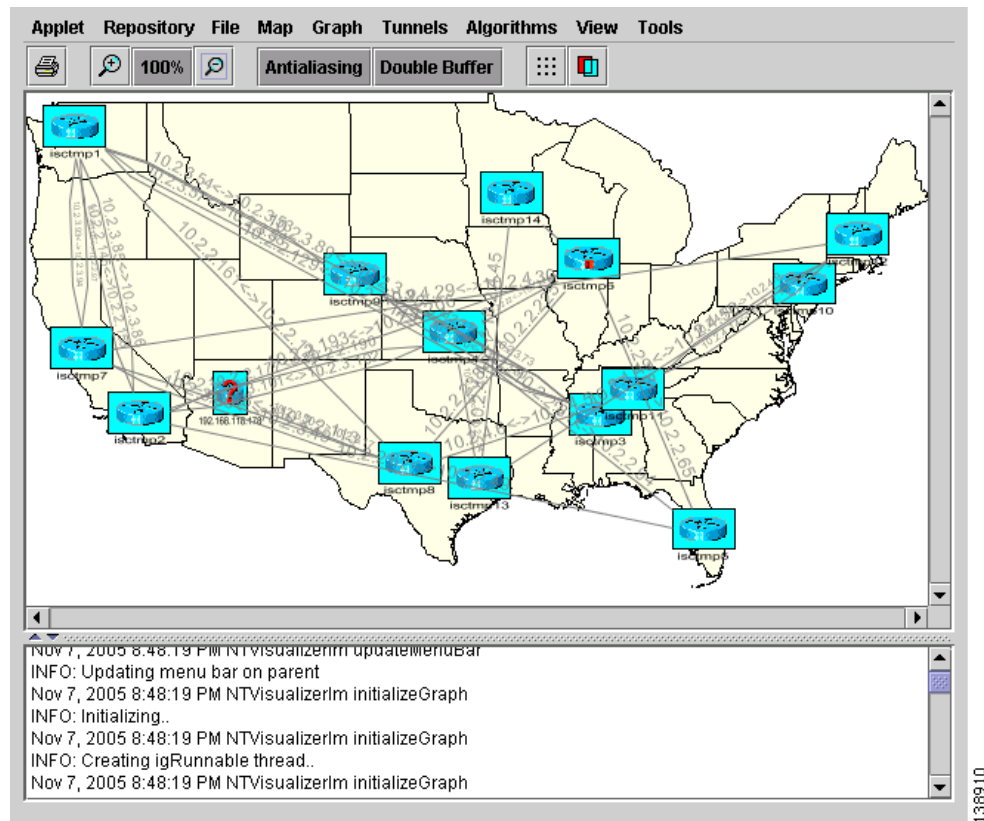
To access the **Map** menu, use the following steps:

-
- Step 1** Choose **Traffic Engineering > TE Topology**.
 - Step 2** Start the **TM Topology Interface Applet**.
If link and node data for your network is already in the repository, a Progress Report lists the various network elements as the corresponding data is loaded.
 - Step 3** Select the **Map** menu.
The menu appears.
From the **Map** menu, you can either load or clear (remove) maps as described in the following.
-

Loading a Map

You might want to set a background map showing the physical locations of the displayed devices. To load a map, use the following steps:

-
- Step 1** In the menu bar, select **Map > Load**.
Providing the web map server is running, the Map Chooser window appears.
 - Step 2** Make your selections in the Map Chooser window.
The right-hand side of the window contains a small control panel, which allows you to select the projection in which a map is shown. A map projection is a projection which maps a sphere onto a plane. Typical projections are Mercator, Lambert, and Stereographic.
For more information on projections, consult the Map Projections section of Eric Weisstein's World of Mathematics at:
<http://mathworld.wolfram.com/topics/MapProjections.html>
If desired, make changes to the settings in the **Longitude Range** and **Latitude Range** fields.
 - Step 3** Select a map file and click **Open** to load the map.
Selecting the map file and clicking the **Open** button starts loading it. Maps can consist of several components and thus a progress dialog is shown informing you which part of the map file is loaded.
A map similar to the one in [Figure 8-28](#) appears.

Figure 8-28 **Loaded Map**

- Step 4** Use the various functions in the menus of the Topology Display window to manipulate the display contents in the Topology view. Some of these are described in subsequent sections.

Adding New Maps

You might need to add your own maps to the selection of maps available to the Topology Tool. This is done by placing a map file in the **\$ISC_HOME/resources/webserver/tomcat/webapps/ipsc-maps/data** directory or a subdirectory thereof within the Prime Provisioning installation. To make this example more accessible, assume that you wish to add a map of Toowong, a suburb of Brisbane, the capital of Queensland. The first step to do so is to obtain maps from a map vendor. All maps must be in the ESRI shape file format (see **ESRI shapefile technical description**). In addition, a data file can accompany each shape file. Data files contain information about objects and the corresponding shapes are contained within the shape file. Let us assume that the vendor provided four files:

- toowong_city.shp
- toowong_city.dbf
- toowong_street.shp
- toowong_street.dbf

We have to create a .map file that informs the TE Topology tool about layers of the map. In this case we have two layers: a city and a street layer. The map file, say, Toowong.map, would thus have the following contents:

```
toowong_city
toowong_street
```

It lists all layers that create a map of Toowong. The order is important, as the first file forms the background layer, with other layers placed on top of the preceding layers.

Having obtained shape and data files and having written the map file, place all five files in the **\$ISC_HOME/resources/webserver/tomcat/webapps/ipsc-maps/data** directory. All map files must be located in this folder. After this is done, the map is automatically accessible to the topology viewer.

Clearing Maps

To clear the active map, select **Map > Clear**.

Use this feature to clear (remove) the active map to leave only nodes and links in the corresponding network.

Using Highlighting and Attributes

The **Graph** menu provides access to a range of tools to manage and manipulate graphs.

Use the JavaServer Pages to look at the list of nodes, links, and tunnels. From the JSP pages, select the display button at the bottom of the window to highlight elements.

The tools in the **Graph** menu serve to modify the appearance of the topology.

These are described in the following sections.

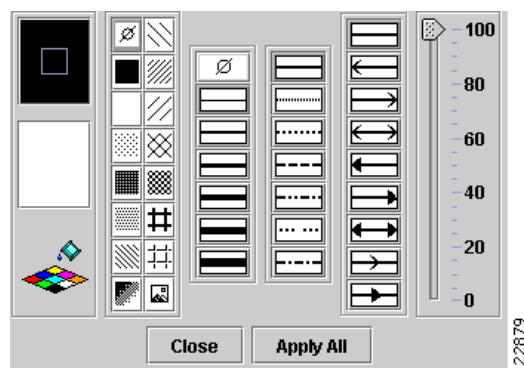
Clear Highlighting

Clear Highlighting serves to remove highlighting from specific elements as listed in its submenus.

Add/Modify Attributes

When you select **Attributes** from the **Graph** menu, the Graphic Attributes window in [Figure 8-29](#) appears.

Figure 8-29 *Graphic Attributes*



The **Add/Modify Attributes** tool is used as follows:

-
- Step 1** Select graph elements (nodes/links) in the topology display.
Use Ctrl/Shift to select multiple elements.
- Step 2** Choose **Graph > Attributes** to open the Graphic Attributes window.
- Step 3** Change the desired attributes and click **Apply All**.



Note Only selected links ([Step 1](#)) are affected.

Clear Current Graph Layout

Use the **Clear** function in the **Graph** menu to remove the topology graph from the current view.

Although this is also achieved with **Layout Graph** in the **Repository** menu, **Layout Graph** re-creates the graph last saved in the repository in addition to clearing the graph.

Using AntiAlias, BackingStore, DoubleBuffer

AntiAlias, found in the **Graph** menu, is used to create smoother lines and a more pleasant appearance at the expense of performance.

BackingStore allows graphics content to be automatically saved when moved to the background and regenerated when returned to the foreground. This helps avoid superfluous refreshing.

DoubleBuffer enables double buffering for dragging elements on the graph.

Using Algorithms

In the **Algorithms** menu various algorithms can be used to enhance and otherwise alter the graph layout.



Note The algorithms only work when the nodes are interconnected with links.

Spring is a graph layout algorithm that optimizes the graph layout based on weights.

Randomize rearranges the nodes in the current topology layout at random.

If there are overlapping links, the layout can be optimized by selecting **Optimize Links**.

The spring settings are used to enhance the appearance of the topology display according to user preferences. When selecting **Spring Settings**, the Spring Settings window appears.

Sample Configlets

The configlets included in this section show the CLIs generated by Prime Provisioning for particular services and features. Each configlet example provides the following information:

- Service

- Feature
- Devices configuration (network role, hardware platform, relationship of the devices and other relevant information)
- Sample configlets for each device in the configuration
- Comments.

All examples in this section assume the presence of an MPLS-TE core.

**Note**

The configlets generated by Prime Provisioning are only the delta between what needs to be provisioned and what currently exists on the device. This means that if a relevant CLI is already on the device, it does not show up in the associated configlet.

This section provides sample configlets for traffic engineering service provisioning in Cisco Prime Provisioning.

It includes the following sections:

- [Primary Tunnel Configlet \(IOS\), page 8-89](#)
- [Bandwidth Protection Backup Tunnel Configlet \(IOS\), page 8-90](#)
- [Connectivity Protection Backup Tunnel Configlet \(IOS\), page 8-91](#)
- [TE Traffic Admission Configlet Using CBTS \(IOS\), page 8-92](#)
- [TE Traffic Admission Configlet \(IOS\), page 8-93](#)
- [Primary Tunnel Configlet \(IOS XR\), page 8-94](#)
- [Bandwidth Protection Backup Tunnel Configlet \(IOS XR\), page 8-95](#)
- [Connectivity Protection Backup Tunnel Configlet \(IOS XR\), page 8-96](#)
- [TE Traffic Admission Configlet Using PBTS \(IOS XR\), page 8-97](#)
- [TE Traffic Admission Configlet \(IOS XR\), page 8-98.](#)

Primary Tunnel Configlet (IOS)

Configuration

- Service: MPLS-TE primary tunnel
- Feature: MPLS TE configlet (IOS) for deploying a primary tunnel
- Device configuration: CISCO12410 with IOS 12.0(32)S.

Configlets

IOS Device Configuration	Comments
<pre> ! Explicit path: ip explicit-path name isctmp2-isctmp8-1 enable next-address 10.2.2.145 next-address 10.2.2.174 ! ! Primary tunnel: interface Tunnel1000 description CISCO ISC-P24 ip unnumbered Loopback0 no ip directed-broadcast tag-switching ip tunnel destination 192.168.118.183 tunnel mode mpls traffic-eng tunnel mpls traffic-eng priority 0 0 tunnel mpls traffic-eng bandwidth 10 tunnel mpls traffic-eng affinity 0x0 mask 0x0 tunnel mpls traffic-eng path-option 1 explicit name isctmp2-isctmp8-1 tunnel mpls traffic-eng path-option 2 dynamic tunnel mpls traffic-eng record-route ! </pre>	<p>Create an explicit path with the specified next addresses, which indicate the strict path that the tunnel traverses.</p> <p>This explicit path is used by the primary tunnel detailed above.</p> <p>Create a TE primary tunnel with the following attributes:</p> <ul style="list-style-type: none"> - tag switching: This command is generated because the policy has the 'mpls ip' flag enabled. This allows the TE tunnels to be used for MPLS VPN traffic. - Destination 192.168.118.183 - TE encapsulation - Setup and hold priorities both 0 - Bandwidth global pool 10 kbps - Tunnel affinity 0x0 - Explicit first path option - Dynamic second path option

Bandwidth Protection Backup Tunnel Configlet (IOS)

Configuration

- Service: MPLS-TE with FRR (Fast Re-Route)
- Feature: This tunnel protects primary tunnel traffic in the event of either a link or node failure
- Device configuration: CISCO12410 with IOS 12.0(32)S.

Configlets

IOS Device Configuration	Comments
<pre>! Explicit path: ip explicit-path name isctmp5-isctmp4-1 enable next-address 10.2.2.145 next-address 10.2.2.174 ! ! Backup tunnel: interface Tunnel1001 description CISCO ISC-B30 ip unnumbered Loopback0 tunnel destination 192.168.118.213 tunnel mode mpls traffic-eng tunnel mpls traffic-eng backup-bw sub-pool 30000 tunnel mpls traffic-eng priority 0 0 tunnel mpls traffic-eng affinity 0x0 mask 0x0 tunnel mpls traffic-eng path-option 1 explicit name isctmp5-isctmp4-1 tunnel mpls traffic-eng record-route ! interface POS0/1 mpls traffic-eng backup-path tunnel 1001 !</pre>	<p>Create an explicit path with the specified next addresses, which indicate the strict path that the tunnel traverses. This explicit path is used by the backup tunnel detailed above.</p> <p>Create a TE backup tunnel with the following attributes:</p> <ul style="list-style-type: none"> - Destination 192.168.118.213 - TE encapsulation - Protect subpool bandwidth of 30000 kbps - Setup and hold priorities both 0 - Tunnel affinity 0x0 - Explicit first path option <p>Backup tunnel 1001 protects interface POS0/1</p>

Connectivity Protection Backup Tunnel Configlet (IOS)

Configuration

- Service: MPLS-TE with FRR (Fast Re-Route)
- Feature: MPLS TE configlet (IOS) for deploying a connectivity protection backup tunnel and its associated exclude address path
- Device configuration: CISCO12410 with IOS 12.0(32)S.

Configlets

IOS Device Configuration	Comments
<pre> ! Explicit path: ip explicit-path name L47-excl enable exclude-address 192.168.1.18 ! ! ! Backup tunnel: interface Tunnel1000 description CISCO ISC-B1 ip unnumbered Loopback0 tunnel mode mpls traffic-eng tunnel destination 10.52.96.38 tunnel mpls traffic-eng priority 0 0 no tunnel mpls traffic-eng bandwidth tunnel mpls traffic-eng path-option 1 explicit name L47-excl tunnel mpls traffic-eng affinity 0x0 mask 0x0 tunnel mpls traffic-eng backup-bw sub-pool unlimited tunnel mpls traffic-eng record-route ! interface ATM4/0.1 point-to-point mpls traffic-eng backup-path Tunnel1000 </pre>	<p>Create an explicit path with an exclude address, which indicates the IP address the path should avoid. This explicit path is used by the backup tunnel detailed above.</p> <p>Create a TE backup tunnel with the following attributes:</p> <ul style="list-style-type: none"> - Destination 10.52.96.38 - TE encapsulation - Setup and hold priorities both 0 - Backup tunnel does not reserve any bandwidth - Explicit first path option - Tunnel affinity 0x0 - Unlimited backup bandwidth for protecting sub pool <p>Set up backup path on ATM interface.</p>

TE Traffic Admission Configlet Using CBTS (IOS)

Configuration

- Service: TE Traffic Admission
- Feature: MPLS TE configlet (IOS) for admitting traffic using Class-Based Tunnel Selection (CBTS)
- Device configuration: CISCO12410 with IOS 12.0(32)S.

Configlets

IOS Device Configuration	Comments
<pre>! TE Traffic Admission using CBTS: interface Tunnel1000 tunnel mpls traffic-eng exp 1 2 3 ! ! Static route: ip route 192.168.118.189 255.255.255.255 Tunnel1000</pre>	<p>Class-based tunnel selection where traffic with EXP bit 1, 2, or 3 are selected</p> <p>Create a static route, which admits all traffic destined for 192.168.118.189 into the above-configured Tunnel 1000.</p>

The above is then deployed to an already existing primary tunnel such as the [Primary Tunnel Configlet \(IOS\)](#), [page 8-89](#).

TE Traffic Admission Configlet (IOS)

Configuration

- Service: TE Traffic Admission
- Feature: MPLS TE configlet (IOS) for TE Traffic Admission
- Device configuration: OSR-7609 with IOS 12.2(33)SRA.

Configlets

IOS Device Configuration	Comments
<pre>! TE Traffic Admission: interface Tunnel1000 tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng autoroute metric relative 0</pre>	Autoroute announce with relative metric, 0 (default)

The above is then deployed to an already existing primary tunnel such as the [Primary Tunnel Configlet \(IOS\)](#), page 8-89.

Primary Tunnel Configlet (IOS XR)

Configuration

- Service: MPLS-TE Primary Tunnel
- Feature: MPLS TE configlet (IOS XR) for deploying a primary tunnel
- Device configuration: CISCO12406 with IOS XR 3.7.0.

Configlets

IOS Device Configuration	Comments
<pre> ! Explicit path: explicit-path name isctmp12-isctmp7-1 index 1 next-address ipv4 unicast 10.163.25.109 index 2 next-address ipv4 unicast 10.163.25.106 ! ! Primary tunnel: interface tunnel-te133 description CISCO ISC-P2 ipv4 unnumbered Loopback0 priority 0 0 signalled-bandwidth 13 destination 192.168.118.214 fast-reroute path-option 1 explicit name isctmp12-isctmp7-1 path-option 2 dynamic record-route ! mpls ldp interface tunnel-te 133 ! </pre>	<p>Create an explicit path with the specified next addresses, which indicate the strict path that the tunnel traverses. This explicit path is used by the primary tunnel detailed above.</p> <p>Create a TE primary tunnel with the following attributes:</p> <ul style="list-style-type: none"> - Destination 192.168.118.214 - TE encapsulation - Setup priority 0 - Hold priority 0 - Reserve 13 kbps from global pool - Tunnel affinity 0x0 - Explicit first path option - Dynamic second path option - Enable FRR for the tunnel <p>Enable ldp (Label Distribution Protocol) on the tunnel interface. This command is generated because the policy has the 'mpls ip' flag enabled. This allows the TE tunnels to be used for MPLS VPN traffic</p>

Bandwidth Protection Backup Tunnel Configlet (IOS XR)

Configuration

- Service: MPLS-TE with FRR (Fast Re-Route)
- Feature: MPLS TE configlet (IOS XR) for deploying a backup tunnel
- Device configuration: CISCO12406 with IOS XR 3.7.0.

Configlets

IOS Device Configuration	Comments
<pre> ! Explicit path: explicit-path name isctmp8-isctmp9-1 index 1 next-address ipv4 unicast 10.163.25.109 index 2 next-address ipv4 unicast 10.163.25.106 ! ! Backup tunnel: interface tunnel-te1009 description CISCO ISC-B1411 ipv4 unnumbered Loopback0 priority 0 0 backup-bw 9600000 destination 10.163.24.131 path-option 1 explicit name isctmp8-isctmp9-1 record-route affinity 0 mask 0 ! mpls traffic-eng interface POS0/1/0/1 backup-path tunnel-te 1009 </pre>	<p>Create an explicit path with the specified next addresses, which indicate the strict path that the tunnel traverses. This explicit path is used by the backup tunnel detailed above.</p> <p>Create a TE backup tunnel with the following attributes:</p> <ul style="list-style-type: none"> - Destination 10.163.24.131 - TE encapsulation - Protect any pool bw of 9600000 kbps - Setup and hold priority of 0 - Tunnel affinity 0x0 - Explicit first path option

Connectivity Protection Backup Tunnel Configlet (IOS XR)

Configuration

- Service: MPLS-TE with FRR (Fast Re-Route)
- Feature: MPLS TE configlet (IOS XR) for deploying a connectivity protection backup tunnel and its associated exclude address path
- Device configuration: CISCO12406 with IOS XR 3.7.0.

Configlets

IOS Device Configuration	Comments
<pre> ! Explicit path: explicit-path name L96-excl index 1 exclude-address ipv4 unicast 192.168.1.42 ! ! ! Backup tunnel: interface tunnel-te1000 description CISCO ISC-B2 ipv4 unnumbered Loopback0 destination 10.52.96.37 priority 0 0 no signalled-bandwidth 0 path-option 1 explicit name L96-excl affinity 0 mask 0 backup-bw sub-pool unlimited record-route ! mpls traffic-eng interface POS0/1/0/2 backup-path tunnel-te 1000 ! </pre>	<p>Create an explicit path with an exclude address, which indicates the IP address the path should avoid. This explicit path is used by the backup tunnel detailed above.</p> <p>Create a TE backup tunnel with the following attributes:</p> <ul style="list-style-type: none"> - Destination 10.52.96.37 - TE encapsulation - Setup priority 0 - Hold priority 0 - Explicit first path option - Tunnel affinity 0x0 - An unlimited sub pool acts as backup bandwidth <p>Tunnel 1000 protects interface POS0/1/0/2</p>

TE Traffic Admission Configlet Using PBTS (IOS XR)

Configuration

- Service: TE Traffic Admission
- Feature: MPLS TE configlet (IOS XR) for admitting traffic using Policy-Based Tunnel Selection (PBTS)
- Device configuration: CISCO12406 with IOS XR 3.7.0.

Configlets

IOS Device Configuration	Comments
<pre>! TE Traffic Admission using PBTS: interface tunnel-tel33 autoroute announce autoroute metric absolute 100 policy-class 2 !</pre>	Autoroute announce with absolute metric 100

The above is then deployed to an already existing primary tunnel such as the [Primary Tunnel Configlet \(IOS XR\)](#), page 8-94.

TE Traffic Admission Configlet (IOS XR)

Configuration

- Service: TE Traffic Admission
- Feature: MPLS TE configlet (IOS XR) for TE Traffic Admission
- Device configuration: CISCO12406 with IOS XR 3.7.0

Configlets

IOS XR Device Configuration	Comments
<pre>! TE Traffic Admission Using Static Route: router static address-family ipv4 unicast 1.2.3.4/32 tunnel-te 1000 123 ! !</pre>	Configuration of TE Traffic Admission on tunnel 1000 with static route

The above is then deployed to an already existing primary tunnel such as the [Primary Tunnel Configlet \(IOS XR\)](#), page 8-94.

Warnings and Violations

This section lists warnings and violations that might be invoked when using the planning tools in Prime Provisioning (computation engine).

Warnings and violations are tied in with the planning tools (see the Planning Tools section in the [Traffic Engineering Management Concepts](#), page 8-111). They are issued under the following circumstances:

- During an attempt to audit, place, repair, or groom a primary managed tunnel.
- During an attempt to protect selected network elements (links, routers, or SRLGs). Here, they help determine the cause of the failed protection (see [Protection Planning](#), page 8-58).

When the off-line backup route generation is called to determine if certain elements can be protected, the backup route generator responds for each element with either a set of tunnels that protect the element or a set of violations and warnings that help determine why the element could not be protected.



Note

In the following, the term DirectedLink refers to a router interface.

This section contains the following:

- [Warnings](#), page 8-99
- [Violations](#), page 8-100

Warnings

This class is characterized by all reports that are warnings. They are considered less severe than violations in the sense that they don't prevent the computation of a protection path.

Protection Computation Warnings

WarningFixVetoed

A fix of this element would have caused a neighbouring element to become unprotected. This fix is vetoed and no changes are proposed.

WarningRouterNotConformant

This element or any adjacent routers is/are not Protocol Conformant. It cannot therefore be protected.

Fields:

- Report Type—Name of report type.
- Description—Description of the problem signaled by the violation.
- Non-conformant router—Router that does not support traffic engineering.

WarningTunnelBandwidthQuotaTooSmall

The bandwidth of a backup tunnel that protects this element is below the minimum allowed bandwidth capacity.

Fields:

- Minimum allowed bandwidth quota—Minimum bandwidth allowed to protect the element in question.
- Actual tunnel bandwidth quota—Actual bandwidth of the backup tunnel.

WarningTunnelNumberTooLarge

There are too many backup tunnels for a flow through this element.

Fields:

- Maximum tunnel number allowed—Maximum number of tunnels allowed for a given network element.
- Actual Tunnel Count—Actual number of tunnels imposed on this network element.
- Flow:
 - Maximum Bandwidth—Maximum bandwidth for the traffic flow that needs to be protected.
 - Head Links—Protected interface for this flow.
 - Through Router —Protected device through which the regular traffic flow passes. If the protected element is a link, the Through Router field will not appear.
 - Tail Router—Hostname of destination (tail) router.
 - Type (NHop, NNHop)—Next hop type: NHOP for link (no through router) and NNHOP for node.

WarningZeroProtectedFlow

A flow through this element is protected by a backup tunnel, but has a maximum flow of zero.

Fields:

- Flow:
 - Maximum Bandwidth—Maximum available bandwidth on the element.
 - Head Links—Protected interface for this flow.
 - Through Router —Protected device through which the regular traffic flow passes. If the protected element is a link, the Through Router field will not appear.
 - Tail Router—Hostname of destination (tail) router.
 - Type (NHop, NNHop)—Next hop type: NHOP for link (no through router) and NNHOP for node.

Violations

This class is specialized by all reports that are violations. They are considered more "severe" than warnings because unlike warnings, they will prevent the computation of a protection path.

Primary Placement Computation Violations

ViolationFrrProtectionInadequate

The FRR protection for a tunnel does not meet the specified protection level.

Fields:

- Report Type—Name of report type.
- Description—Description of the problem signaled by the violation.
- Required FRR Protection Level—Used to enable an MPLS traffic engineering tunnel to use a backup tunnel in the event of a link failure if a backup tunnel exists. Possible levels are **None**, **Best Effort**, **Link and SRLG**, and **Link, SRLG and Node**.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
- Path—Tunnel Path
 - Node—Device hostname. Is only displayed if the protection level is "Link, SRLG & Node".
 - Protected (Node)—Indicates whether each node is protected (Yes) or not (No). Is only displayed if the protection level is "Link, SRLG & Node".
 - Link Label—IP addresses of the interfaces on the link.
 - Protected (Link)—Indicates whether each link is protected (Yes) or not (No).

ViolationInconsistentResourceAttributeChanges

A Topology-change attempts to modify one or more attributes on a resource causing a pair of its attributes to become inconsistent.

Fields:

- Report Type—Quality report, warning report, or violation report.

- Description—Description of the problem signaled by the violation.
- Resource—
 - Id—Id for head device or head interface representing the network resource.
 - Type—Resource device or interface.
- Attributes:
 - Attribute—Names of inconsistent attributes.
 - New Value—New attribute value proposed by user.

ViolationInconsistentTunnelAttributeChanges

A Tunnel-change attempts to modify one or more attributes on a tunnel causing a pair of its attributes to become inconsistent.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
- Attributes:
 - Attribute—Names of inconsistent attributes.
 - New Value—New attribute value proposed by user.

ViolationLinkAffinityMismatch

A least one directed link in the path of a Primary Tunnel does not have attribute flags that match the affinity bits and mask of the Tunnel.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Affinity Bits/Mask—Affinity bits and mask of the tunnel.
- Path—Name of tunnel path.
 - Outgoing Interface—Hostname/IP address of outgoing interface.
 - Attribute Flags—Links attributes to be compared to the tunnel's affinity bits. All have to be identical to have a valid path. The violation is triggered when at least one is different.

ViolationLinkPoolOversubscribed

The specified bandwidth pool for a directed link is over-subscribed by Primary Tunnels that pass through it.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Directed Link:
 - Head Device/Interface—Hostname for the head device and IP address of interface.
 - Tail Device/Interface—Hostname for the destination (tail) device or interface.
 - Pool—Global pool or sub pool.
 - Pool Bandwidth—The allocated global pool or sub pool bandwidth on the link.
- Primary Tunnel (table)—Specifies how many tunnels are using the link resource.
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Bandwidth—Total bandwidth of the tunnel.
 - Pool—Global pool or sub pool.
 - Path—Name of tunnel path.

ViolationMaxReRoutesExceeded

This number of Primary Tunnel re-routes in this solution exceeds the specified maximum.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Number of re-routes in solution—Number of re-routes proposed by the computation engine.
- Specified maximum number of re-routes—Maximum number of re-routes allowed.

ViolationNoPathInLayout

In the presence of other Primary Tunnels that have already been placed on the topology, no legitimate path is possible for a requested Primary Tunnel. Note: If a user requested path was specified this only means that the Primary Tunnel could not be placed on that requested path in the presence of other Primary Tunnels.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Requested Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Bandwidth—Total bandwidth of the tunnel.
 - Requested Path—User-specified path for the tunnel.
 - Pool—Global pool or sub pool.

- FrrProtection—Possible protection levels are **None**, **Best Effort**, **Link and SRLG**, and **Link, SRLG and Node**.
- Propagation Delay—The time it takes for traffic to travel along a link from the head interface to the tail interface.
- AffinityBits/Mask—Affinity bits and mask of the tunnel.

ViolationNoPathInTopology

Irrespective of other Primary Tunnels placed upon the topology, no valid path is possible for a requested Primary Tunnel. Note: If a user requested path was specified this only means that the Primary Tunnel could not be placed on that requested path irrespective of other tunnels.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Requested Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of (destination) tail router.
 - Bandwidth—Total bandwidth of the tunnel.
 - Requested Path—User-specified path for the tunnel.
 - Pool—Global pool or sub pool.
 - FrrProtection—Possible protection levels are **None**, **Best Effort**, **Link and SRLG**, and **Link, SRLG and Node**.
 - Propagation Delay (optional)—The maximum time allowed for traffic to travel along the requested path.
 - AffinityBits/Mask—Affinity bits and mask of the tunnel.

ViolationNoTunnelForDemand

No path implements a requested PrimaryTunnel, even though there exists a valid path in the network that this tunnel could take.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Requested Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Bandwidth—Total bandwidth of the tunnel.
 - Requested Path—User-specified path for the tunnel.
 - Pool—Global pool or sub pool.
 - FrrProtection—Possible protection levels are **None**, **Best Effort**, **Link and SRLG**, and **Link, SRLG and Node**.

- Propagation Delay (optional)—The maximum time allowed for traffic to travel along the requested path.
- AffinityBits/Mask—Affinity bits and mask of the tunnel.

ViolationPathMismatch

A Primary Tunnel has a different path to that specified for it in the User Specified Path.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Actual Path—Actual path of the tunnel associated with the violation.
 - Requested Path—User-specified path for the tunnel.

ViolationPathNotConnected

The path of a Primary Tunnel is not “connected”, that is, it does not form a connected sequence of admin-up links between the tunnel head and tail, or it contains loops.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Path—Name of tunnel path.

ViolationPathUsesMissingLinks

A Tunnel-change attempts to create or modify a Tunnel so that its path or “User Requested Path” uses one or more directed links that do not exist in this topology.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Change Type—Add Tunnel/Modify Tunnel.
 - Path Type—Requested/Actual.

- Path—Name of tunnel path.
- Outgoing Interface—Yes or No depending on whether a link is missing.
- Incoming Interface—Yes or No depending on whether a link is missing.

ViolationPrimaryTunnelDelayTooLong

A Primary Tunnel has a propagation delay that is larger than the Maximum Propagation Delay specified for it.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Required Max Propagation Delay—The maximum time allowed for traffic to travel along the requested path.
- Primary Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.
 - Path—Name of tunnel path.
 - Actual Propagation Delay (table)—The time it takes for traffic to travel along each link in the entire path.
 - Link—Link segments in path.
 - Propagation Delay—Travel time for the traffic for each link segment.

ViolationResourceIdUnknown

A change attempts to remove or modify a resource (link, router or SRLG) with an Id, when no resource with that Id exists.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Resource to be removed:
 - Id—Id for head device or head interface representing the network resource.
 - Type—Resource device or interface.

ViolationTunnelIdInUse

A change attempts to add a Primary Tunnel with an Id that already exists.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Tunnel to Add:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.

- Existing Tunnel:
 - Name—Tunnel identifier composed of a name and a tunnel number.
 - Head—Hostname of head router.
 - Tail—Hostname of destination (tail) router.

ViolationTunnelIdUnknown

A change attempts to remove or modify a Primary Tunnel with an Id when no tunnel with that Id exists.

Fields:

- Report Type—Quality report, warning report, or violation report.
- Description—Description of the problem signaled by the violation.
- Tunnel to Remove:
 - Id—Unique tunnel identifier used within Prime Provisioning.

Protection Computation Violations

ViolationAggregateBandwidthOnLink

The bandwidth of backup tunnels for this element, which pass through the link, have a maximum bandwidth quota that exceeds the backup bandwidth of the link.

Fields:

- Required Bandwidth (due to tunnels)—Required bandwidth for the tunnels on the link.
- Link:
 - Backup Bandwidth—Total available bandwidth of the link.
 - Head Router—Hostname of the head router.
 - Head Interface—IP address of the head interface.
 - Tail Router—Hostname of destination (tail) router.
 - Tail Interface—IP address of the destination (tail) interface.
 - Label—IP addresses of the interfaces on the link.
 - Admin Status—Indicates whether the link is **Up** or **Down**.

ViolationBadBackupTunnel

The tunnel does not protect a flow over this element.

ViolationBandwidthProtectionMismatch

The tunnel backup bandwidth quotas of all the tunnels protecting a flow do not add up exactly to the maximum bandwidth of that flow.

Fields:

- Protected bandwidth—The protectable bandwidth of the protection path.
- Flow:
 - Maximum Bandwidth—Maximum available bandwidth on the element.
 - Head Links—Protected interface for this flow.

- Through Router —Protected device through which the regular traffic flow passes. If the protected element is a link, the Through Router field will not appear.
- Tail Router—Hostname of destination (tail) router.
- Type (NHop, NNHop)—Next hop type: NHOP for link (no through router) and NNHOP for node.

ViolationLinkLevelTunnelDelayTooLarge

The delay of the backup tunnel is greater than that allowed.

Fields:

- Maximum allowed delay—Maximum delay allowed on the backup tunnel.
- Actual delay of tunnel—Actual delay of the backup tunnel.

ViolationNoBackupTunnels

There are no backup tunnels protecting this flow through the element.

Fields:

- Flow:
 - Maximum Bandwidth—Maximum available bandwidth on the element.
 - Head Links—Protected interface for this flow.
 - Through Router —Protected device through which the regular traffic flow passes. If the protected element is a link, the Through Router field will not appear.
 - Tail Router—Hostname of destination (tail) router.
 - Type (NHop, NNHop)—Next hop type: NHOP for link (no through router) and NNHOP for node.

ViolationPassesThroughSRLG

A backup tunnel is protecting a flow over this element that starts at a link within an Shared risk link group(SRLG). However that tunnel also passes through another link in the same SRLG.

Fields:

- Link:
 - Backup Bandwidth—Total available bandwidth of the link.
 - Head Router—Hostname of the head router.
 - Head Interface—IP address of the head interface.
 - Tail Router—Hostname of destination (tail) router.
 - Tail Interface—IP address of the destination (tail) interface.
 - Label—IP addresses of the interfaces on the link.
 - Admin Status—Indicates whether the link is **Up** or **Down**.
- SRLG—User-defined SRLG name.
- Flow:
 - Maximum Bandwidth—Maximum available bandwidth on the element.
 - Head Links—Protected interface for this flow.

- Through Router —Protected device through which the regular traffic flow passes. If the protected element is a link, the Through Router field will not appear.
- Tail Router—Hostname of destination (tail) router.
- Type (NHop, NNHop)—Next hop type: NHOP for link (no through router) and NNHOP for node.

ViolationUsesFailedElement

A backup tunnel that protects this element also uses it.

Document Type Definition (DTD) File

The Document Type Definition (DTD) file provides the rules required by the XML import file for importing bulk data into Prime Provisioning.

For instructions on how to import tunnels into Prime Provisioning, see [Import Primary Tunnel, page 8-49](#).

This section includes the following:

- [DTD File, page 8-108](#)
- [Example, page 8-111](#)

DTD File

This is the DTD file provided with Prime Provisioning.

```
<?xml version="1.0" encoding="UTF-8"?>

<!-- Data Definition for file based tunnel import -->

<!-- Import File Structure -->
<!ELEMENT IMPORT_DATA (TUN_ADD|TUN_CHANGE|TUN_DELETE|TUN_MIGRATE)+ >

<!-- Notes on attributes:
importId: must be unique within the file,
         it is alphanumeric, must begin with alpha character,
         and no special character
head, tail: hostname of valid TE enabled device
policy: name of existing managed tunnel policy
bw: must be numeric and values between 0-2147483647
tnum: is the number portion of a tunnel interface
      E.g. for "interface tunnel3", use tnum="3"
      must be numeric and values between 0-65535
-->

<!-- Tunnel Add

- #IMPLIED attributes are optional, if not specified, defaults to null
- If tnum is not specified, system will generate tunnel number
- To enable auto bandwidth, specify AUTOBW element
- bw is required if autobw is not enabled
- By default, tunnel will be created with a system path and a dynamic path
```

```
-->

<!ELEMENT TUN_ADD (AUTOBW?)>
<!ATTLIST TUN_ADD
    importId ID #REQUIRED
    head CDATA #REQUIRED
    tail CDATA #REQUIRED
    policy CDATA #REQUIRED
    bw CDATA #IMPLIED
    tnum CDATA #IMPLIED>

<!-- Tunnel Change

    - #IMPLIED attributes are optional, if not specified, value on existing
      tunnel is kept
    - To enable auto-bw, or to change auto-bw parameters, specify AUTOBW element
    - To disable auto-bw, set disableAutoBw="yes" and do not specify AUTOBW element
    - Existing tunnel path cannot be changed directly, setting reroutable="true"
      will enable system to reroute the tunnel if necessary

-->

<!ELEMENT TUN_CHANGE (AUTOBW?)>
<!ATTLIST TUN_CHANGE
    importId ID #REQUIRED
    head CDATA #REQUIRED
    tnum CDATA #REQUIRED
    policy CDATA #IMPLIED
    bw CDATA #IMPLIED
    disableAutoBw (yes) #IMPLIED
    reroutable (true|false) #IMPLIED>

<!-- Tunnel Delete

    - all attributes are required to identify tunnel to be deleted

-->

<!ELEMENT TUN_DELETE EMPTY>
<!ATTLIST TUN_DELETE
    importId ID #REQUIRED
    head CDATA #REQUIRED
    tnum CDATA #REQUIRED>

<!-- Tunnel Migrate

    - #IMPLIED attributes are optional, if not specified, value on existing
      tunnel is kept
    - All comments under Tunnel Change (above) applies to Tunnel Migrate
    - only unmanaged primary tunnel can be migrated
    - for tunnels with unmanaged tunnel policy, must specify a managed policy
    - for tunnels that was non-conformant:
      . if bw was zero, specify a new bw or enable auto-bw
      . if path was dynamic or non-conformant, the path options will be
        replaced with a system path and a dynamic path, and reroutable will
        be set to true.
    - reroutable attribute applicable only for tunnel that had a conformant first
      explicit path (i.e. explicit path with no loopback)

-->
```

```

<!ELEMENT TUN_MIGRATE (AUTOBW?)>
<!ATTLIST TUN_MIGRATE
    importId ID #REQUIRED
    head CDATA #REQUIRED
    tnum CDATA #REQUIRED
    policy CDATA #IMPLIED
    bw CDATA #IMPLIED
    disableAutoBw (yes) #IMPLIED
    reroutable (true|false) #IMPLIED>

<!-- Auto Bandwidth

- #IMPLIED attributes are optional, if not specified, value is set to null
  for TUN_ADD and existing value is kept TUN_CHANGE
- maxBw is required when used in TUN_ADD or if existing tunnel is not auto-bw
  enabled
- minBw and maxBw must be numeric and values between 0-2147483647
- maxBw must be greater than minBw if specified
- freq must be numeric and values between 300-604800

-->

<!ELEMENT AUTOBW EMPTY>
<!ATTLIST AUTOBW
    freq CDATA #IMPLIED
    minBw CDATA #IMPLIED
    maxBw CDATA #IMPLIED>
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE IMPORT_DATA SYSTEM "TeImport.dtd">

<IMPORT_DATA>

<!-- Add New Managed Tunnel -->
<TUN_ADD importId="a1" head="isctmp3" tail="isctmp1" policy="mgdPolicy" bw="400" />
<TUN_ADD importId="a2" head="isctmp2" tail="isctmp9" policy="mgdPolicy" >
    <AUTOBW freq="300" minBw="100" maxBw="200"/>
</TUN_ADD>

<!-- Modify Existing Tunnel -->
<TUN_CHANGE importId="c1" head="isctmp2" tnum="200" bw="30" />
<TUN_CHANGE importId="c2" head="isctmp4" tnum="2" policy="mgdPolicy" reroutable="true"/>
<TUN_CHANGE importId="c3" head="isctmp5" tnum="46">
    <AUTOBW freq="300" minBw="100" maxBw="200"/>
</TUN_CHANGE>
<TUN_CHANGE importId="c4" head="isctmp2" tnum="200" bw="30" disableAutoBw="yes"/>

<!-- Delete Existing Tunnel -->
<TUN_DELETE importId="d1" head="isctmp3" tnum="45"/>

<!-- Migrate Tunnel -->
<TUN_MIGRATE importId="m1" head="isctmp2" tnum="3" policy="mgdPolicy"/>
<TUN_MIGRATE importId="m2" head="isctmp5" tnum="1" policy="mgdPolicy"/>

</IMPORT_DATA>

```

Example

The following is an example of a tunnel import XML file conforming to the DTD file specified in [DTD File, page 8-108](#). It consists of a sample block for each of the Add, Change, Delete, and Migrate operations.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE IMPORT_DATA SYSTEM "TeImport.dtd">

<IMPORT_DATA>

<!-- Add New Managed Tunnel -->
<TUN_ADD importId="a1" head="isctmp3" tail="isctmp1" policy="mgdPolicy" bw="400" />
<TUN_ADD importId="a2" head="isctmp2" tail="isctmp9" policy="mgdPolicy" >
  <AUTOBW freq="300" minBw="100" maxBw="200"/>
</TUN_ADD>

<!-- Modify Existing Tunnel -->
<TUN_CHANGE importId="c1" head="isctmp2" tnum="200" bw="30" />
<TUN_CHANGE importId="c2" head="isctmp4" tnum="2" policy="mgdPolicy" reroutable="true"/>
<TUN_CHANGE importId="c3" head="isctmp5" tnum="46">
  <AUTOBW freq="300" minBw="100" maxBw="200"/>
</TUN_CHANGE>
<TUN_CHANGE importId="c4" head="isctmp2" tnum="200" bw="30" disableAutoBw="yes"/>

<!-- Delete Existing Tunnel -->
<TUN_DELETE importId="d1" head="isctmp3" tnum="45"/>

<!-- Migrate Tunnel -->
<TUN_MIGRATE importId="m1" head="isctmp2" tnum="3" policy="mgdPolicy"/>
<TUN_MIGRATE importId="m2" head="isctmp5" tnum="1" policy="mgdPolicy"/>

</IMPORT_DATA>
```

Traffic Engineering Management Concepts

This chapter includes an overview of Cisco Prime Provisioning and of some of the concepts used in this guide. This chapter includes the following sections:

- [Prime Provisioning TEM Overview, page 8-112](#)
- [Features in Prime Provisioning, page 8-112](#)
- [Prime Provisioning TEM Basics, page 8-112](#)
 - [Managed/Unmanaged Primary Tunnels, page 8-112](#)
 - [Conformant/Non-Conformant Tunnels, page 8-113](#)
 - [Multiple Concurrent Users, page 8-114](#)
 - [Multiple OSPF Areas, page 8-115](#)
 - [Bandwidth Pools, page 8-116](#)
 - [Planning Tools, page 8-117](#)
 - [Connectivity Protection \(CSPF\) Backup Tunnels, page 8-118](#)
 - [Class-Based Tunnel Selection, page 8-118](#)

- [Policy-Based Tunnel Selection, page 8-119.](#)

Prime Provisioning TEM Overview

TEM is the Traffic Engineering Management module of Prime Provisioning. It is a tool for managing Multiprotocol Label Switching Traffic Engineering (MPLS TE) primary tunnels and backup tunnels for the purpose of offering traffic Service Level Agreement (SLA) guarantees. It provides bandwidth protection management, network discovery, and support for configuring MPLS TE. It includes a number of powerful planning tools, including a sophisticated primary path calculation tool and backup tunnel calculation for element protection.

MPLS TE mechanisms are provided to support requirements for predictability, traffic flow matched to QoS requirements, and Fast Restoration with Guaranteed Bandwidth, ensuring that strict SLA performance criteria (availability, delay, jitter) are met.

Features in Prime Provisioning

Prime Provisioning adds a range of MPLS TE primary tunnel management features:

- Tunnel Audit—finding inconsistencies after making tunnel modifications
- Tunnel Admission—admitting new tunnels onto the network
- Tunnel Repair—fixing tunnel inconsistencies after network and service changes
- Network Grooming—optimizing global network utilization.

In addition, Prime Provisioning offers interaction and integration with Prime Provisioning features:

- Service activation focus
- Integration with other Prime Provisioning modules
- Data Persistence
- Logging of user intent
- Service state management
- Service auditing
- Web-based GUI
- Role-Based Access Control (RBAC).

Prime Provisioning TEM Basics

To understand how Prime Provisioning works, you need to first know certain key concepts.

Managed/Unmanaged Primary Tunnels

In Prime Provisioning, the concept of managed tunnels is at the center of TE planning activities.

It is important to understand the differences:

- Managed TE tunnels:
 - (setup/hold) priority zero

- non-zero RSVP bandwidth
- explicit first path option
- auto bandwidth must have a max value
- Unmanaged tunnels: All other tunnels.

In the Prime Provisioning Graphical User Interface (GUI), there is a separate entry point for dealing with managed and unmanaged tunnels.

Conformant/Non-Conformant Tunnels

Understanding the concepts of conformant and non-conformant tunnels is key to making the most efficient use of Prime Provisioning.

Prime Provisioning only allows the creation of conformant tunnels. Non-conformant tunnels can be introduced through the TE Discovery process (see [TE Network Discovery](#), page 8-10 of the User Guide).

Defining Conformant/Non-Conformant Tunnels

In the Prime Provisioning design, a sharp distinction has been made between conformant and non-conformant tunnels:

- **Conformant tunnel**—A well-behaved tunnel that meets Prime Provisioning's TE management paradigm (described below). A managed tunnel can only be a conformant tunnel. A non-zero priority unmanaged tunnel would also be a conformant tunnel. However, a conformant tunnel is not necessarily a managed tunnel.

A connectivity protection tunnel is marked Conformant = true if it has zero tunnel bandwidth, unlimited backup bandwidth, and an 'exclude address' first path option. For the BW Protected setting, a tunnel should have a defined non-zero backup bandwidth, and a strict path option 1.

- **Non-conformant tunnel**—A TE tunnel, which might impact Prime Provisioning's ability to meet bandwidth guarantees. This could be due to unknown bandwidth requirements such as no max bandwidth configured for auto-bandwidth, potential for pre-emption, dynamic paths, etc. A zero priority unmanaged tunnel would also be a non-conformant tunnel.

The following are examples of non-conformant tunnels:

- a tunnel with zero setup and hold priority, an explicit first path option, but with zero bandwidth;
- a tunnel with zero setup and hold priority, a non zero bandwidth, but with a dynamic first path option;
- a tunnel with zero setup and hold priority, an explicit path option of 1 and an auto bandwidth without a maximum defined.;
- a connectivity protection tunnel marked Conformant = false is reserved for backup tunnels, which have neither zero tunnel bandwidth, unlimited backup bandwidth, or an 'exclude address' first path option.

Why are the above tunnels non-conformant? Because Prime Provisioning attempts to manage all tunnels with zero setup and hold priority, to ensure the links they pass through all have sufficient bandwidth, are affinity consistent, and do not break delay or FRR constraints defined in the TE policy.

But if the tunnel's path is dynamic or the amount of bandwidth it requires is undefined, Prime Provisioning does not have the information with which to manage the tunnel, so it marks it as non-conformant. All the non-conformant tunnels are displayed in the TE Unmanaged Primary Tunnels SR window.

Managing Non-Conformant Tunnels

It is important to understand that non-conformant tunnels not only might cause the SLAs to be violated, they might also have an adverse effect on the managed tunnels (taking away bandwidth from them, for example).

However, when a non-conformant tunnel is discovered, a warning is logged. Prime Provisioning tracks non-conformant tunnels so that they can be decommissioned.

So conformant tunnels are preferred. They allow the system to offer bandwidth guarantees for managed tunnels. Unmanaged non-conformant tunnels might or might not provide the needed bandwidth and no bandwidth guarantees are given.

The action to take when you have non-conformant tunnels is either to change the setup and hold priorities to non-zero values (so they cannot preempt the managed tunnels) or migrate them to managed tunnels, allowing the tool to find a suitable explicit path.

Multiple Concurrent Users

In previous releases TEM only supported a single GUI user. This release introduces support for multiple concurrent users, for all browsing, updating, and provisioning operations.

Concurrent Use with Managed and Unmanaged Tunnels

To understand how the multiple user feature is implemented in TEM, it is important to understand the difference between a managed and an unmanaged tunnel. This is described in the section [Managed/Unmanaged Primary Tunnels, page 8-112](#).

There are important differences between how managed and unmanaged tunnels are handled when it comes to multiple user support:

- For managed tunnels, an SR encapsulates all managed tunnels. A SR operation might optimize all the objects within the snapshot following path computations performed by the Router Generator server.
- For unmanaged tunnels, an SR is defined as a tunnel-head end router. Thus, with unmanaged tunnels there are certain restrictions. For example, two users cannot concurrently provision on the same device.
- TEM prevents Unmanaged Tunnel SRs from provisioning concurrently on the same device but supports Unmanaged Tunnel SRs provisioning concurrently on different devices.
- All managed tunnels are contained within a shared Managed TE Tunnel SR for each TE Provider. For unmanaged tunnels, a distinct Unmanaged TE Tunnel Service Request is created per head device. TEM supports multiple SRs per TE Provider.

Multiple TEM users can browse and provision in TEM. Up to 20 concurrent users are supported, of which up to seven can perform provisioning tasks.

Previously all primary tunnels, managed and unmanaged were in a single TE tunnel SR per TE provider. Now, to facilitate multiple simultaneous changes to managed tunnels, the TE Tunnel SR has been split into one managed tunnel SR per TE provider and one unmanaged tunnel SR per head TE router.

Parallel provisioning is not possible on the same SR, but because SRs exist at router level for unmanaged tunnels, unmanaged tunnels can be provisioned on separate routers at the same time.

Locking Mechanism

When an unmanaged tunnel is provisioned, the head TE router of the tunnel is locked. This can be seen on the TE Nodes window in the System Lock Status column. The locking prevents any other user from deploying any kind of tunnel to that router until the provisioning task completes and the TE router is unlocked.

The locking mechanism also applies to other Prime Provisioning features, such as backup tunnels, resource SRs, link deletion, and TE traffic admission. Resource SRs include deleting/editing explicit paths, deleting protected elements, deleting/editing SRLG's, etc.

In the case of link deletion, a level of intelligence is built in. When there are no more tunnels to be rerouted or deleted by the user or Prime Provisioning and left with TE associated objects alone, a user intervention will be required to carry out link deletion. As part of this deletion, if there are any backup tunnels protecting any interfaces that have been selected for deletion, the locking mechanism will be in place during deployment of backup tunnels. For further information about deleting TE links, see [Deleting TE Links, page 8-24](#) of the User Guide.

Some of the potential errors you might encounter are described in [Locking Operation Errors, page 8-78](#) of the User Guide.

When a managed primary tunnel or a backup tunnel is provisioned, the TE provider it is associated with is locked. This can be seen on the TE Provider window in the System Lock Status column. A lock at the TE provider level prevents another user from making any tunnel change on this TE provider, irrespective of which TE router the tunnel starts at.

The reason why the locking mechanism of managed tunnels and backup tunnels is different from that of unmanaged tunnel is that the managed tunnels and backup tunnels use a path generation algorithm to find an optimal route for the tunnel that fulfills all constraints, and this algorithm needs a stable global view of the TE topology and all the tunnels in it on which to base its routing decisions. This can only be achieved by allowing only one user to make changes at one time.

For more information about how to manage Prime Provisioning locking mechanism, see [Managing the Locking Mechanism, page 8-77](#) of the User Guide.

Multiple OSPF Areas

Prime Provisioning supports the discovery, management, and provisioning of TE Tunnels within multiple Open Shortest Path First (OSPF) areas.

Prime Provisioning only manages primary and backup TE tunnels within the scope of an OSPF area. There is no support for the discovery and creation of inter-OSPF areas.

In Prime Provisioning, an OSPF area is represented by a TE provider. After an area is assigned to a TE provider, it might not be changed. Multiple TE providers can be associated with one Prime Provisioning provider.

Devices Suitable for TE Discovery

In a network with multiple OSPF areas, where each OSPF area is represented by a TE provider, any router in an OSPF area can be used for TE Discovery. Using multiple TE providers (multiple OSPF areas) under one provider allows the provisioning of inter-area L3VPN.

**Note**

Prime Provisioning will not discover or provision inter-area TE tunnels (those with a head router in one area and a tail router in a different area).

To discover a multi area network, you have to discover each area in turn using TE Discovery (see [TE Network Discovery, page 8-10](#) of the User Guide). The seed node can be any device within that area, including an Area Border Router (ABR).

TE Discovery and the TE Area Identifier

TE Discovery is associated with a TE provider and each TE provider is assigned an area. The area is assigned during the process of creating the TE Provider (see [Creating a TE Provider, page 8-7](#) of the User Guide) and can be a simple integer value or dotted decimal notation, Area 0.6.0.0 for example.

TE provider objects are aware of which area they are responsible for, either specified on creation or automatically populated during discovery, and will accommodate conversion between Dot notation and Decimal notation, defaulting to the notation used in the network.


When discovery is run against an area with a selected TE provider, all tunnels and explicit paths associated with that area will be imported into the Prime Provisioning database. The steps for performing a per area discovery are documented in the [Managing Per Area Discovery, page 8-15](#) of the User Guide.

Example of Multiple OSPF Area Network

TE routers within a TE provider can be assigned to different regions, for example on a geographical basis, so that devices are grouped in regions in a logical way. Also, Prime Provisioning allows you to filter by region. Assigning objects to specific regions is a manual task that is carried out after discovery from **Inventory > Provider Devices**. Here the region of any PE device can be changed via the Select Region pop-up window.

In the following example [Figure 8-30](#), two TE providers are each responsible for one OSPF area that is created and visualized under one Prime Provisioning provider.

Figure 8-30 Multiple OSPF Areas Network Diagram

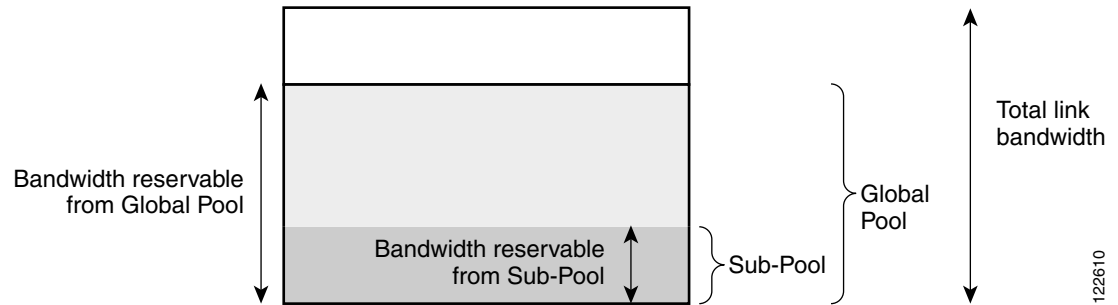


For instructions on how to manage TE providers, see [Creating a TE Provider, page 8-7](#) of the User Guide.

Bandwidth Pools

The bandwidth of each TE enabled interface is assigned a number of nested bandwidth pools. Currently, IOS supports two, namely Global Pool and Sub Pool.

For a better understanding of bandwidth pools, see [Figure 8-31](#).

Figure 8-31 Bandwidth Pools

As [Figure 8-31](#) illustrates, Sub Pool is nested inside Global Pool. Thus, if a primary tunnel reserves bandwidth from the Sub Pool, it will also reserve the same bandwidth from the Global Pool.

Bandwidth reservations (primary tunnels) from the Sub Pool must not exceed, in total, the Sub Pool size. Likewise, bandwidth reservations from the Global Pool must not exceed, in total, the Global Pool size.

Planning Tools

They are intended for evaluating planned improvements to a traffic-engineered network based on What-If scenarios.

The planning tools include the following features:

- Primary planning tools:
 - Tunnel Audit—Audits for inconsistencies in primary placement on the existing network with or without proposed tunnel or resource changes.
 - Tunnel Placement—Usually for new tunnels. Tunnel Placement can generate a new route. It can be used for a tunnel that did not have a path before and needs to be placed.
 - Tunnel Repair—Logically performed after Tunnel Audit (if something is wrong). Tunnel Repair has rerouting capabilities and can be used to move tunnels.
 - Grooming—An optimization tool that works on the whole network. It is only available when no tunnel attributes have been changed.
- Protection planning tools:
 - Audit SR—Audits protection for manually added, modified, and deleted backup tunnels before they are deployed.
 - Compute Backup—Automatically calculates the optimal backup tunnel for selected network elements.
 - Audit Protection—Audits protection of the selected elements against the existing backup tunnels.

The planning tools are fully integrated within Prime Provisioning and are available from various locations within the GUI:

- TE Protected Elements (Compute Backup and Audit Protection)
- Create Managed TE Tunnel (Tunnel Audit, Tunnel Placement, Tunnel Repair, Grooming)
- Create TE Backup Tunnel (Audit SR).

Connectivity Protection (CSPF) Backup Tunnels

In addition to the bandwidth-protected backup tunnels created by TEM, you can create a set of CSPF-routed backup tunnels within Prime Provisioning. These CSPF-routed backup tunnels are managed from the TE Protection SR window.

A connectivity protection backup tunnel uses an “exclude-address” explicit path. This explicit path is created in the TE Explicit Path List window. An exclude address path is different from a strict path in that instead of listing the hops the path should use, it lists the hops the path should avoid. The CSPF algorithm on the router will make the decision as to which precise path to use, but it will be constrained to not be able to use the hops in the exclude address path configuration. This sort of path is particularly useful for backup tunnels, as the interfaces the exclude address path should avoid can be the interfaces that the backup tunnel is protecting.

In Prime Provisioning, these backup tunnels are configured with unlimited backup bandwidth. Unlimited means no bandwidth is guaranteed, but as much as is available at the time of the failure will be used. So in effect the bandwidth protection is best effort but the connectivity is guaranteed. Connectivity protection backup tunnels can be used in addition to or instead of bandwidth protection backup tunnels.

Differences between bandwidth protection and connectivity protection backup tunnels:

- A bandwidth protection backup tunnel has a strict explicit path as its first path option, whilst a connectivity protection tunnel has an exclude address explicit path as its first path option.
- A bandwidth protection backup tunnel has a defined backup bandwidth whilst a connectivity tunnel has unlimited backup bandwidth on a best effort basis.
- A bandwidth protection backup tunnel is passed to the Route Generator algorithm which generates optimal backup tunnels and verifies existing tunnels fully protect the elements, whereas connectivity protection tunnels are not passed to the algorithm and it is up to you to ensure they are fulfilling their purpose.

Class-Based Tunnel Selection

Multi-Protocol Label Switching Traffic Engineering Class-Based Tunnel Selection (CBTS) enables you to dynamically route and forward traffic with different class of service (CoS) values onto different TE tunnels between the same tunnel head end and the same tail end. The packet's CoS values are located in the EXP bits. There are 8 EXP bits, numbered 0 to 7.

The set of TE (or DS-TE) tunnels from the same head end to the same tail end can be configured to carry different CoS values. After configuration, CBTS dynamically routes and forwards each packet into the tunnel that:

- is selected for traffic admission using the standard autoroute or static route mechanisms, and
- has EXP bits matching that of the packet.

Thus CBTS is not a form of traffic admission to TE tunnels directly, it is rather an additional criteria that traffic must satisfy before being admitted to tunnels via the autoroute or static route mechanisms that TEM supports.

Because CBTS offers dynamic routing over DS-TE tunnels and requires minimum configuration, it greatly eases deployment of DS-TE in large-scale networks. CBTS can distribute all CoS values onto many different tunnels.

The CBTS feature has the following restrictions:

- For a given destination, all CoS values are carried in tunnels terminating at the same tail end. Either all CoS values are carried in tunnels or no values are carried in tunnels. In other words, for a given destination, you cannot map some CoS values in a DS-TE tunnel and other CoS values in a Shortest Path First (SPF) Label Distribution Protocol (LDP) or SPF IP path.
- CBTS does not allow load-balancing of a given EXP value in multiple tunnels. If two or more tunnels are configured to carry a given experimental (EXP) value, CBTS picks one of these tunnels to carry this EXP value.
- The operation of CBTS is not supported with Any Transport over MPLS (AToM), MPLS TE Automesh, or label-controlled (LC)-ATM.

When traffic admission to tunnels is achieved using global static routes, and when there is more than one tunnel to a given destination with the same administrative weight, the CBTS attribute acts as a tiebreaker in selecting the right tunnel. (See above discussion of load-balancing with CBTS.)

Policy-Based Tunnel Selection

Multi-Protocol Label Switching Traffic Engineering Policy-Based Tunnel Selection (PBTS) enables you to dynamically route and forward traffic based on a policy onto different TE tunnels between the same tunnel head end and the same tail end. The routing algorithm is performed on the headend router's ingress interface prior to forwarding lookup.

In the Prime Provisioning implementation of PBTS, traffic is directed into specific TE tunnels using the interface command `policy-class`. Whereas CBTS is aimed at IOS devices, PBTS is strictly designed for IOS XR devices.

Like CBTS, PBTS is not a form of traffic admission to TE tunnels directly, but rather an additional criteria that traffic must satisfy before being admitted to tunnels via the autoroute or static route mechanisms that TEM supports.



Note

Prime Provisioning itself does not provision the policy class, it merely associates a tunnel with an existing policy class. This is done by specifying the policy-class attribute in the range 1 to 7.

For more information on CBTS, see [Class-Based Tunnel Selection](#), page 8-118.

For general information on PBTS and IOS XR, see

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/mpls/configuration/guide/gc37te.html#wp1325561.

