



CHAPTER 4

Traffic Engineering Management Concepts

This chapter includes an overview of Cisco Prime Fulfillment and of some of the concepts used in this guide. This chapter includes the following sections:

- [Prime Fulfillment TEM Overview, page 4-1](#)
- [Features in Prime Fulfillment, page 4-2](#)
- [Prime Fulfillment TEM Basics, page 4-2](#)
 - [Managed/Unmanaged Primary Tunnels, page 4-2](#)
 - [Conformant/Non-Conformant Tunnels, page 4-3](#)
 - [Multiple Concurrent Users, page 4-4](#)
 - [Multiple OSPF Areas, page 4-5](#)
 - [Bandwidth Pools, page 4-6](#)
 - [Planning Tools, page 4-7](#)
 - [Connectivity Protection \(CSPF\) Backup Tunnels, page 4-8](#)
 - [Class-Based Tunnel Selection, page 4-8](#)
 - [Policy-Based Tunnel Selection, page 4-9.](#)

Prime Fulfillment TEM Overview

TEM is the Traffic Engineering Management module of Prime Fulfillment. 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 Fulfillment

Prime Fulfillment 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 Fulfillment offers interaction and integration with Prime Fulfillment features:

- Service activation focus
- Integration with other Prime Fulfillment modules
- Data Persistence
- Logging of user intent
- Service state management
- Service auditing
- Web-based GUI
- Role-Based Access Control (RBAC).

Prime Fulfillment TEM Basics

To understand how Prime Fulfillment works, you need to first know certain key concepts.

Managed/Unmanaged Primary Tunnels

In Prime Fulfillment, 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 Fulfillment 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 Fulfillment.

Prime Fulfillment only allows the creation of conformant tunnels. Non-conformant tunnels can be introduced through the TE Discovery process (see [Chapter 36, “TE Network Discovery”](#) of the User Guide).

Defining Conformant/Non-Conformant Tunnels

In the Prime Fulfillment design, a sharp distinction has been made between conformant and non-conformant tunnels:

- Conformant tunnel—A well-behaved tunnel that meets Prime Fulfillment’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 Fulfillment’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 Fulfillment 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 Fulfillment 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 Fulfillment 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 4-2](#).

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 Fulfillment 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 Fulfillment 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 37-5](#) of the User Guide.

Some of the potential errors you might encounter are described in [Locking Operation Errors, page 42-10](#) 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 Fulfillment locking mechanism, see [Managing the Locking Mechanism, page 42-9](#) of the User Guide.

Multiple OSPF Areas

Prime Fulfillment supports the discovery, management, and provisioning of TE Tunnels within multiple Open Shortest Path First (OSPF) areas.

Prime Fulfillment 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 Fulfillment, 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 Fulfillment 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 Fulfillment 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 [Chapter 36, “TE Network Discovery”](#) 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 35-8](#) 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 Fulfillment database. The steps for performing a per area discovery are documented in the [Managing Per Area Discovery, page 36-5](#) 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 Fulfillment 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 4-1](#), two TE providers are each responsible for one OSPF area that is created and visualized under one Prime Fulfillment provider.

Figure 4-1 Multiple OSPF Areas Network Diagram




For instructions on how to manage TE providers, see [Creating a TE Provider, page 35-8](#) 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 4-2](#).

Figure 4-2 Bandwidth Pools



As [Figure 4-2](#) 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 Fulfillment 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 Fulfillment. 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 Fulfillment, 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 Fulfillment 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 Fulfillment 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 4-8](#).

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.

