WHITE PAPER

CISCO IP SOLUTION CENTER TRAFFIC-ENGINEERING MANAGEMENT

An application in the Cisco® IP Solution Center (ISC)—A family of network-intelligent element-management applications for planning, provisioning, and troubleshooting Multiprotocol Label Switching (MPLS) and Metropolitan (Metro) Ethernet networks

SUMMARY

MPLS has rapidly become an important technology for provisioning and managing core networks. One of the main reasons for this is its ability to support network convergence, allowing delay- and jitter-sensitive traffic (for example, voice) to be carried over the cost-effective infrastructure of IP networks. MPLS Traffic Engineering (MPLS-TE) provides a means to ensure the bandwidth and delay characteristics of the network and to address connectivity and bandwidth protection in failure scenarios. It also enables major improvements to network usage by providing a mechanism for avoiding congestion in one portion of the network and under-usage in another.

Unlocking the value of MPLS-TE requires the support of a traffic-engineering management system. Determining an optimal path for a given tunnel may be dependent upon many other tunnels and MPLS-TE configurations in the network; managing multiple potential conflicting constraint requirements such as bandwidth and delay, generating a globally optimized layout of tunnels, and generating an efficient bandwidth protection solution using MPLS-TE Fast Reroute (FRR) technology all require sophisticated routing algorithms. Cisco IP/MPLS Traffic Engineering Management incorporates world-class hybrid optimization techniques to provide major usage improvements. Cisco IP/MPLS Traffic Engineering Management:

- Generates the paths for tunnels that meet constraints, including bandwidth, DiffServ-Aware Traffic Engineering (DS-TE) pool, affinity, delay, and protection level
- Automatically discovers and audits the tunnels on the network
- Manages, optimizes, and deploys tunnels on the network

Cisco IP Solution Center Traffic Engineering Management (ISC:TEM) incorporates a carrier-class provisioning engine that removes the potential for configuration errors due to manual processes. In the MPLS-TE environment this is a critical component because configuration activities can be lengthy and detailed. Cisco IP Solution Center Traffic Engineering Management incorporates a graphical- and table-based display of MPLS-TE, helping the user quickly search through information, visualize the data, and associate tunnels and resources on the network.

CHALLENGES

Supporting the Converged Network

To use MPLS-TE effectively in support of network optimization, high-value service-level agreement (SLA) traffic or bandwidth protection requires MPLS-TE path management. Path management is part of the configuration management function and includes the calculation of primary paths for MPLS-TE tunnels, re-optimization of paths, and the calculation of protection paths. To support traffic on the converged network, traffic-engineering tunnel paths must be calculated with reference to multiple requirements such as bandwidth, delay, and network protection. As an example, voice traffic may have a maximum network delay of 70 ms and must cross only links that are protected against failure.

Constraint Shortest Path First (CSPF) is the default mechanism for determining paths through the network, and this provides a fast solution for basic path calculations that involve only one constraint. This solution cannot provide true network-wide optimization, support for multiple constraints, or the efficient calculation of network protection.

Protecting the Network

As voice and data networks converge, the necessary service qualities for each type of traffic (Frame Relay, ATM, IP, or time-division multiplexing (TDM) have to be supported by a single network, which must be able to offer high levels of availability and resilience, including rapid recovery times and bandwidth guarantees if a network failure occurs. Historically, such services have been provided by building redundancy that is used only if a failure occurs—such schemes include standard SONET/SDH protection.

Supporting Voice over IP

Operators are beginning to exploit the potential of voice over IP (VoIP) as a way of offering traditional voice services more cheaply then using the traditional public switched telephone network (PSTN) solution. Additionally, VoIP offers a solution for rapidly providing customers new and innovative services that provide additional sources of revenue.

The challenge of supporting VoIP is to provide the same end-user experience that voice customers expect from the PSTN over what has traditionally being seen as a best-effort service. In particular, availability, delay, and jitter all must meet certain requirements if users are to accept the VoIP solution.

SOLUTION

Cisco IP Solution Center Traffic Engineering Management is a centralized traffic-engineering planning and provisioning tool that takes advantage of Cisco IOS® Software MPLS-TE to provide the advanced features specified previously. Cisco ISC:TEM is a Web services-based application with a server and repository running on a Solaris system with GUI access through a Web browser. The Cisco ISC framework offers a highly scalable operations-support-system (OSS) solution with a distributed architecture allowing processes to be load balanced across different physical systems.

Network Discovery

Network discovery is performed by using a set of command-line interface (CLI) commands that discover the MPLS-TE topology and any existing tunnels on the network. Cisco ISC:TEM connects to the traffic engineering-enabled routers using either Telnet or Secure Shell (SSH) Protocol, and a seed device is defined for network discovery. During the discovery process any existing tunnels are checked for consistency and then saved in the repository along with node and link information. Cisco ISC:TEM then considers the repository to contain the correct view of the traffic-engineering network. Subsequent discoveries detect any differences between the network and the repository and raise a discrepancy report that alerts the operator to any differences. It is then the operator’s responsibility to determine what is correct: the network or the repository.

The discovery process relies on the flooding of traffic-engineering attributes through the network as defined in RFCs 3630 and 3784 and, therefore, operates in a single Open Shortest Path First (OSPF) area or Intermediate System-to-Intermediate System (IS-IS) level. Figure 1 illustrates the interaction of Cisco ISC:TEM with the network for both discovery and provisioning. The discovery process is as follows:

- Traffic-engineering extensions to OSPF and IS-IS flood traffic-engineering topology and tunnel information around the network.
- Cisco ISC:TEM communicates with a predefined seed router to discover the traffic-engineering network topology.
- Each traffic-engineering router that has been discovered is queried for traffic-engineering tunnel and interface information.
- Discovered tunnels are audited for consistency and then stored in the repository with the topology information.

Figure 1. Cisco ISC:TEM Network Discovery and Provisioning

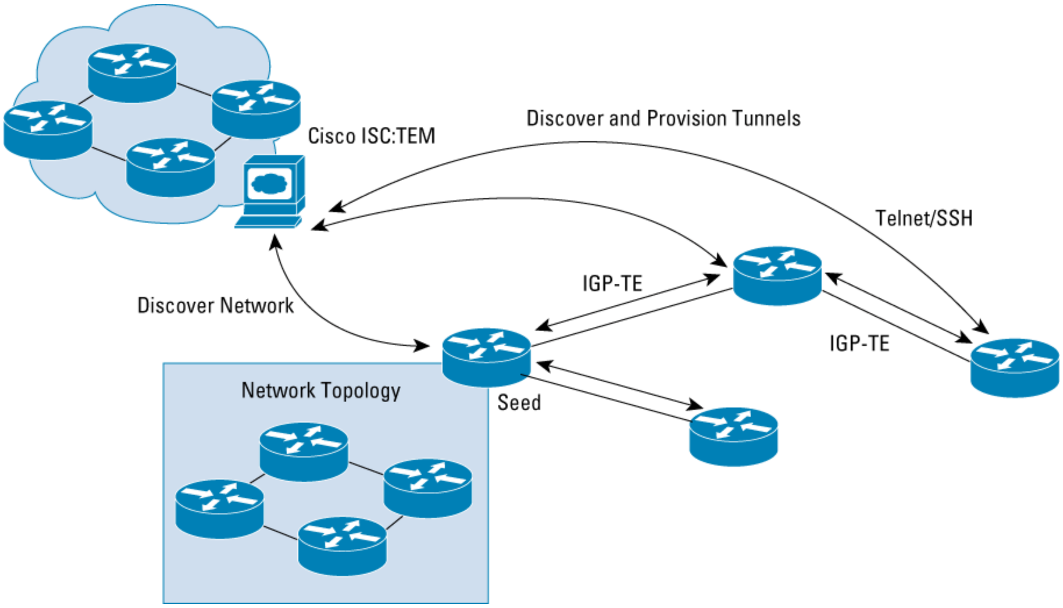

When discovering existing tunnels, Cisco ISC:TEM expects tunnels to have the following set of criteria in order to become a managed tunnel:

- Nonzero tunnel bandwidth or maximum auto bandwidth value configured
- Explicit first-path option
- Zero setup and hold priorities

If these three conditions are met, the tunnel is managed by Cisco ISC:TEM. Other tunnels are not managed by Cisco ISC:TEM until they are “admitted” into Cisco ISC:TEM, meaning that Cisco ISC:TEM calculates a new primary path for the tunnel.

Network Protection

Traditionally operators have turned to SONET/SDH protection to provide fast restoration from network faults. MPLS FRR offers the same fast protection but implemented in Layer 3. FRR offers connectivity protection with a failover time in the order of 50 ms, and Cisco ISC:TEM extends this to enable bandwidth protection as well. Links, nodes, and shared-risk link groups (SRLGs) can be protected using Cisco ISC:TEM, which calculates the tunnels required and then provisions them on the network.

These tunnels are then activated by the routers adjacent to the network-element failure, and any primary tunnels crossing the protected element are placed into these backup tunnels. It is important to note that it is the traffic-engineering bandwidth through an element that is being protected, not the primary tunnels. This means that after the protection tunnels are calculated and provisioned to protect a network element, any changes to primary tunnels crossing that element do not result in changes to the protection tunnels. Additionally, the backup tunnels can protect many primary tunnels through the use of MPLS label stacking. Therefore, this protection solution is very scalable because changes are infrequent and the number of protection tunnels depends only on the number of elements being protected.

Figure 2 shows how the bandwidth protection algorithm works. Each link is divided into a primary pool and a backup pool. The primary pool is the traffic-engineering bandwidth pool, which can be the global pool or the subpool, and when the algorithm is calculating protection, it is this primary pool that must be protected. The remaining bandwidth on a link is considered available for carrying the backup traffic, although this can be modified in Cisco ISC:TEM. If a failure occurs, the operator can be confident that the normal traffic on the link (that is, the primary pool) and the protected traffic from the element that failed do not exceed the link capacity.

Figure 2. Bandwidth Protection Pools

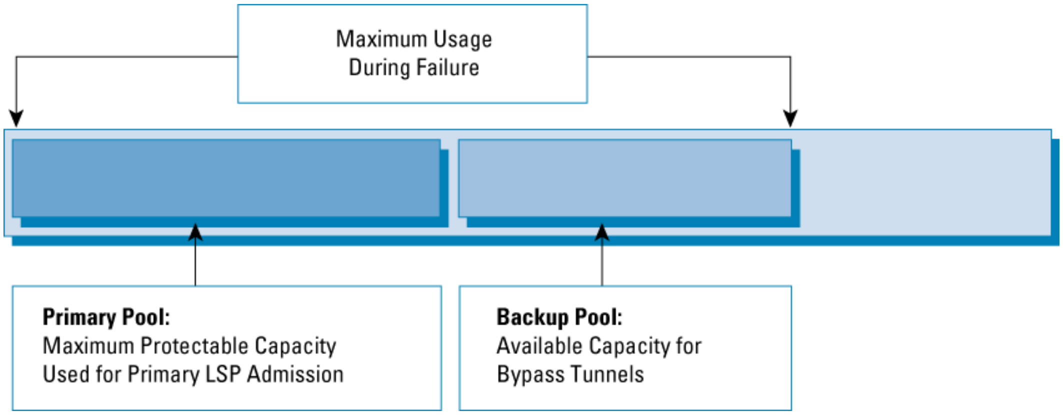

When calculating backup tunnels, Cisco ISC:TEM uses patented techniques to minimize the amount of backup bandwidth that is required to protect a network. This combined with the bandwidth sharing between independent element failure cases means that there is more remaining bandwidth left to carry revenue-producing traffic. An example of how the algorithm can minimize the amount of backup bandwidth required is shown in Figure 3.

This example shows the amount of bandwidth that must be protected on each link if the central node fails. The algorithm tries to protect all possible traffic flows between the adjacent routers that pass through the failed router *regardless of the status of any existing primary tunnels*. At first glance it would appear that there are 15 units of backup bandwidth to be protected from R5 through R1 to R4 and 15 units between R2 to R4 though R1. But the link between R1 and R4 can carry at most only 20 units of backup bandwidth, meaning only 20 units of backup bandwidth must be found instead of 30.

Figure 3. Backup Bandwidth Sharing

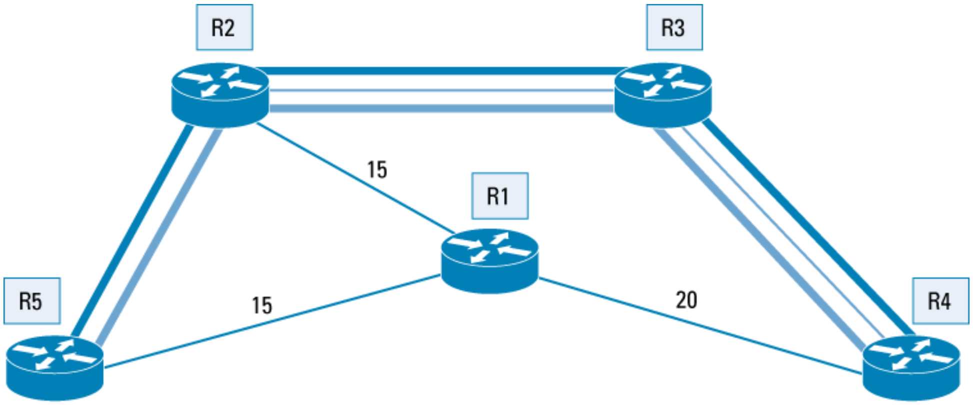

Through this and other techniques, Cisco ISC:TEM can offer a much more efficient backup solution than either CSPF or SONET/SDH, and this can be seen in Figure 4. It shows the amount of backup bandwidth required in the network to support 1 Gbps of primary traffic. For SONET linear automatic protection switching (APS), an extra 1 Gbps of bandwidth is required to protect each 1 Gbps of traffic because of the 1 + 1 protection mechanism. The best CSPF can do is the same as SONET, whereas Cisco ISC:TEM needs at most an extra 800 Mbps of bandwidth. This produces savings in terms of the amount of bandwidth that cannot carry revenue-producing traffic. If Cisco ISC:TEM is used, there is more bandwidth available to carry this revenue traffic.

Figure 4. Extra Bandwidth Required to Protect Links in Network

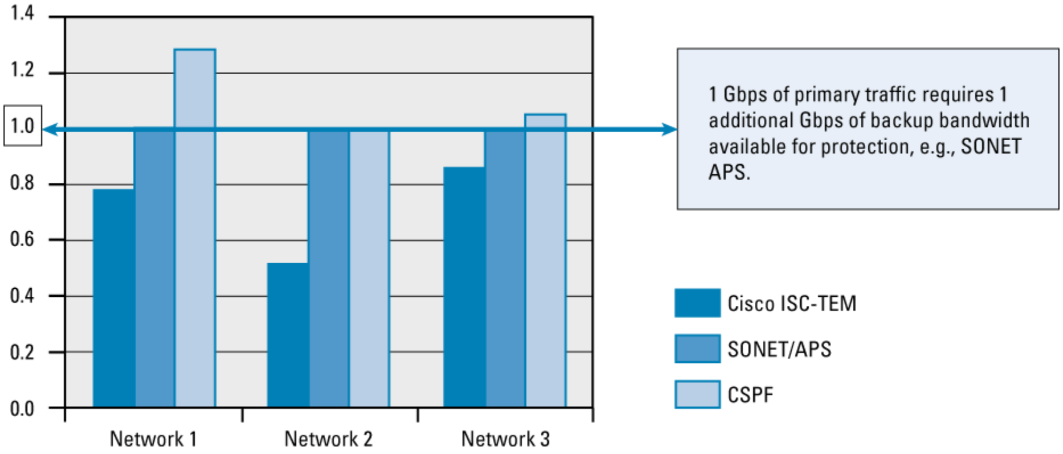

If Cisco ISC:TEM is used to calculate backup tunnels and to keep track of the bandwidth being protected, the backup tunnels can be signaled on the network with “0 bandwidth”, meaning that the backup tunnels do not reserve any network resources. This powerful technique allows sharing of backup bandwidth between failure cases and also allows low-priority traffic to use the backup bandwidth when it is not in use.

Tunnel Path Placement

Cisco ISC:TEM calculates the placement of tunnels on the network to support the desired traffic requirements. Traffic demands are entered by the user or loaded through an Extensible Markup Language (XML) file. A user can specify the bandwidth required for the demand along with any delay or protection constraints. Protection requirements are specified by the type of protected element that must be traversed by the demand. Protection requirement are listed below:

- Link + SRLG protected
- Link + SRLG + Node protected
- Protection is not a concern.

After the user specifies a set of new demands, they are passed to the route-generator algorithm, which computes paths for the new tunnels based on the demand requirements, the network topology, and the existing tunnels. By default, any existing tunnels are not moved during a tunnel placement request but as with most Cisco ISC:TEM functions, this can be modified to suit the operator’s environment. The algorithm performs the following steps when calculating paths through the network:

- Maximize the amount of new bandwidth being requested on the network.
- Minimize the worst-case subpool use.
- Minimize the worst-case global-pool use.
- Minimize the traffic-engineering metrics.
- During these steps the demand constraints must always be met; for example, delay constraint.

After the calculation is complete, the results are returned in a table for reviewing, and on the graphical display the user can see the paths the tunnels will take through the network. These can then be provisioned onto the network and traffic can be admitted onto the tunnels. It is important to note that when the tunnels are provisioned they do not carry traffic until the operator admits traffic onto the tunnels. The operator enables traffic admission through Cisco ISC:TEM and can schedule when this occurs, so that the change can be implemented during a change window.

Cisco ISC:TEM offers the potential for significant savings when compared to other methods by allowing a network to carry more traffic and defer upgrades as traffic grows. Figure 5 shows the difference between CSPF and Cisco ISC:TEM when traffic on a network increases and the impact of this on network capacity is investigated. As the network traffic increases, links become saturated and must be upgraded. In the example, when the use of a link reaches 70 percent the capacity is doubled. The graph shows the number of links that must be upgraded when using Cisco ISC:TEM compared with CSPF. There is a clear cost saving when using Cisco ISC:TEM over CSPF, with only one link requiring upgrading compared to a minimum of seven links when using CSPF. This reduction in the number of links offers clear return-on-investment (ROI) benefits in terms of line rental and router interface-card costs.

Figure 5. Network Upgrade Scenarios for Cisco ISC:TEM and CSPF-Based Solutions

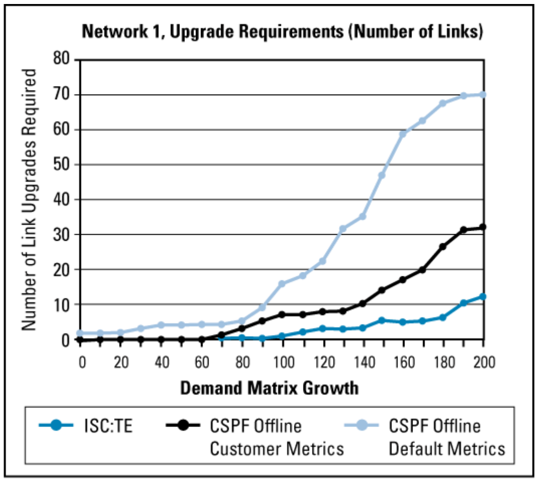

Grow Network Traffic and Upgrade
Links to Keep Up Sufficient Capacity

Required Network Upgrades for each algorithm:	100% Demand Growth		200% Demand Growth	
	No. of Links	B/W Gbps	No. of Links	B/W Gbps
ISC:TEM	1	2	12	7
CSPF Offline, Customer Metrics	7	5	32	21
CSPF Offline, Default Metrics	1	13	70	58

Case Study Notes	
Tier 1 International Network	
Nodes, links (approx.) =	175, 380
Upgrade a link when primary utilization >	70%
When upgrading, multiply link capacity by	2
Use-edge-to-edge as constraint?	No

Cisco ISC:TEM can offer the improvements in performance over CSPF because of the nature of the algorithm as well as the techniques used to derive a solution. In summary, the Cisco ISC:TEM algorithms are search-based algorithms and can account for all the tunnels simultaneously. CSPF, on the other hand, is a “greedy” algorithm that places tunnels one at a time with no reference to any of the other tunnels being placed.

Planning and Implementing Network Changes

Cisco ISC:TEM allows the impact of network changes on MPLS-TE tunnels to be investigated before the change goes ahead. For example, the impact of taking a link out of service to upgrade the line cards can be investigated, and any tunnels that are affected can be rerouted. Similarly, if the delay figure for a link needs adjusting, the impact on the end-to-end delay of any delay constrained tunnels crossing that link can be investigated before the delay figures are adjusted. When a change is proposed with Cisco ISC:TEM, it runs a tunnel audit, which checks to see if any primary tunnels are directly affected by the change and if the change will affect the protection of any other tunnels in the network.

If tunnels are affected by a change, then a tunnel fix is run to reroute those affected tunnels. The steps the algorithm takes are as follows:

- Assume all tunnels in the network can be moved to accommodate the change. (The user can set one or more tunnels to be unmovable.)
- Reroute the affected tunnels and move any other tunnels if required.
- Minimize the number of tunnels that are rerouted to accommodate the change.
- The tunnel constraints must be held during this process.

The operator can then decide if the solution is appropriate or not and can proceed with the change or cancel it. Through this feature, Cisco ISC:TEM offers a “what-if” planning capability that allows the operator to see the effect of various changes on the network before actually implementing them.

Network Grooming

If changes are made to the network and the tunnels crossing the network, then the network usage may become suboptimal. Cisco ISC:TEM offers a solution to this through the network grooming function. This function simply takes all the existing tunnels on the network and recalculates their paths to optimize the network usage. Following are the steps the algorithm follows:

- Assume all tunnels can be moved. (The user can override this.)
- Minimize the worst-case subpool usage.
- Minimize the worst-case global-pool usage.
- Minimize the traffic-engineering metrics.

An example of when the network grooming function can be used is when a large number of new tunnels have been admitted to the network over a period of time. In the typical (also default) operation of the tunnel-placement algorithm, existing tunnels on the network are not moved when new tunnels are added. Therefore, over time this can result in a suboptimal use of the network and a network groom can be used bring the network back to an optimal state.

Network Provisioning

Network provisioning, also performed through Telnet or SSH connections, uses the Cisco ISC provisioning engine to provide a robust, carrier-class mechanism for configuring traffic-engineering tunnels. During the provisioning process the routers are queried for their configuration, which is then compared with the new configuration. This results in a set of commands called a configlet, which is sent to the router. The final step of the provisioning process is to fetch the router configuration again and audit the results against the Cisco ISC:TEM view of what the configuration should look like. It is only after this auditing step has successfully completed that a tunnel is considered to be deployed. The commands that were sent to the router also are stored in the repository and can be reviewed at any time.

The same process is used when changes are being made to the network attributes, for example, bandwidth pool sizes, and these occur after any required tunnel moves have been completed. When the user decides to remove a link from the service, the tunnels are moved in anticipation of the outage, but the link is not disabled by Cisco ISC:TEM. This must be done by the user in case other traffic is crossing the link, for example, Interior Gateway Protocol (IGP) traffic.

After deployment, the tunnels can be audited at any time and the status of the tunnel along with the path it is following through the network can be compared with what is in the repository.

Integration with Existing OSS

Cisco ISC:TEM provides an XML-based interface to add new demands or make changes to existing demands. The interface provides an alternative to the GUI for making bulk updates, but the user must manually initiate the placement calculations. Tunnels are specified and then the import button is used to check the consistency of the data in the file and then import into Cisco ISC:TEM. In this way Cisco ISC:TEM can be integrated into the service-provisioning cycle while allowing the operator to react to the result of the tunnel calculation results (refer to Figure 6).

Figure 6. Importing Tunnel Configuration Information from File

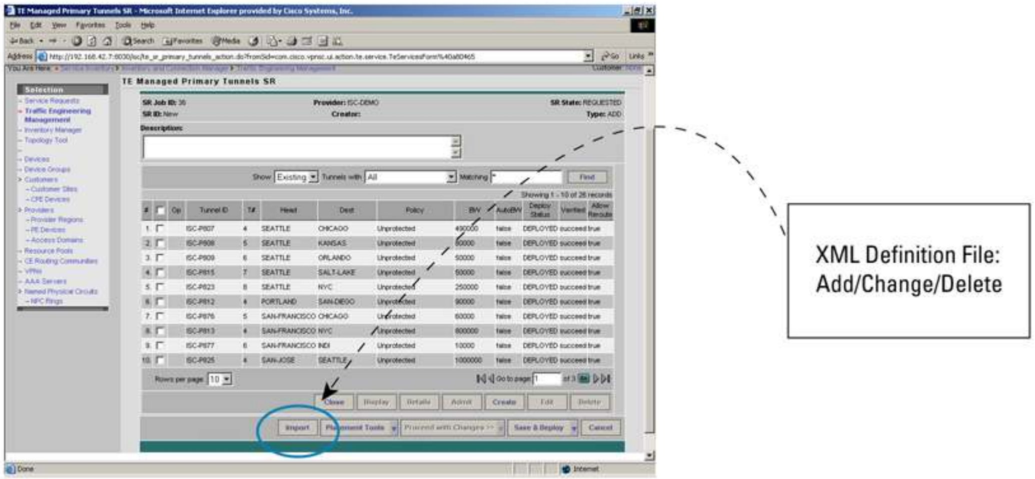

TE Managed Primary Tunnels SR

SR Job ID: 36 Provider: ISC-DEMO SR State: REQUESTED
SR ID: New Creator: Type: ADD

Description:

Show: Existing Tunnels with: All Matching: *

Showing 1 - 10 of 26 records

#	☐	Op	Tunnel ID	T#	Head	Dest	Policy	BW	AutoBW	Deploy Status	Verified	Allow Reroute
1.	☐		ISC-P807	4	SEATTLE	CHICAGO	Unprotected	490000	false	DEPLOYED	succeed	true
2.	☐		ISC-P808	5	SEATTLE	KANSAS	Unprotected	80000	false	DEPLOYED	succeed	true
3.	☐		ISC-P809	6	SEATTLE	ORLANDO	Unprotected	50000	false	DEPLOYED	succeed	true
4.	☐		ISC-P815	7	SEATTLE	SALT-LAKE	Unprotected	50000	false	DEPLOYED	succeed	true
5.	☐		ISC-P823	8	SEATTLE	NYC	Unprotected	250000	false	DEPLOYED	succeed	true
6.	☐		ISC-P812	4	PORTLAND	SAN-DIEGO	Unprotected	90000	false	DEPLOYED	succeed	true
7.	☐		ISC-P876	5	SAN-FRANCISCO	CHICAGO	Unprotected	60000	false	DEPLOYED	succeed	true
8.	☐		ISC-P813	4	SAN-FRANCISCO	NYC	Unprotected	800000	false	DEPLOYED	succeed	true
9.	☐		ISC-P877	6	SAN-FRANCISCO	INDI	Unprotected	10000	false	DEPLOYED	succeed	true
10.	☐		ISC-P825	4	SAN-JOSE	SEATTLE	Unprotected	1000000	false	DEPLOYED	succeed	true

Rows per page: 10 Go to page: 1 of 3

XML Definition File:
Add/Change/Delete

SUMMARY

MPLS-TE provides the tools for operators to support converged networks offering Layer 2 and 3 services, VoIP, and fast restoration. Cisco IP Solution Center Traffic Engineering Management is an effective management solution enabling operational efficiency and network optimization for service providers and enterprises using MPLS-TE.

REFERENCES:

- Bandwidth Protection for Fast Reroute—White paper: http://www.cisco.com/go/mpls/
- RFC 3630—Traffic Engineering Extensions to OSPF
- RFC 3784—Intermediate System-to-Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204105.58_ETMG_LF_12.04

Printed in the USA