White Paper

Deploying Guaranteed-Bandwidth Services with MPLS

Guaranteed-Bandwidth Services with MPLS

Introduction

Customers who are building a high-speed IP core can enable new revenue and cost savings by offering additional services using Multiprotocol Label Switching (MPLS) technology. A Service Provider can take advantage of MPLS technology, to offer guaranteed-bandwidth services. With guaranteed bandwidth solutions, customers now can offer voice and data services with point-to-point guarantees in an MPLS network. Guaranteed bandwidth service enables Service Provider to offer predictable packet delivery characteristics at various network conditions and loads.

This solutions document describes how a hypothetical MPLS network can deliver guaranteed-bandwidth services using traffic engineering and quality of service (QoS). It discusses various implementation trade-offs and includes a preferred network design that outlines “best practices” for actual guaranteed-bandwidth services deployment. The guaranteed-bandwidth solutions presented in this document are based on recommended customer configurations. These solutions were tested and verified in a lab environment and can be deployed in the field. Alternative ways to implement guaranteed-bandwidth solutions are not discussed in this guide.

This solutions document describes basic design and deployment of a network capable of delivering bandwidth guarantees on an IP-based platform. It does not discuss in detail the general operation of the protocols associated with deployment, such as Resource Reservation Protocol (RSVP), Label Distribution Protocol (LDP), and differentiated services (DiffServ), nor does it discuss the management and automation aspect for service provisioning or deploying QoS features.

This document contains the following sections:

- Customer Business Objectives
- Proposed Solution: Guaranteed Bandwidth Services with MPLS
- Implementation of Proposed Solution: Guaranteed Bandwidth with MPLS
- Related Documents

Customer Business Objectives

Customer business objectives in deploying guaranteed-bandwidth services include:

- Take advantage of existing network infrastructure to create additional revenue streams and provide value-added services to customers by enabling point-to-point guarantees and connectivity.
- Enable predictable packet-delivery characteristics between any two nodes on the network for real-time applications, such as high-quality video and audio, voice-over-IP (VoIP) traffic, and distance learning, for content providers and subscribers.

Guaranteed-Bandwidth Services Definition

Today's multiservice packet networks rely on IP-based packet switching. However, IP networks built without QoS simply offer a best-effort service, and do not provide strict delay, jitter, or bandwidth guarantees required for VoIP and other real-time traffic. Using Cisco IOS® Software QoS and the Internet Engineering Task Force (IETF) DiffServ model for QoS, the network treats VoIP traffic appropriately.

Although bandwidth is fairly inexpensive today, even in networks with ample bandwidth an “insurance policy” is essential to ensure guaranteed quality for voice, regardless of the overall network traffic load. Thus, a service must extract the maximum profit benefit from every bit of bandwidth available. Although the DiffServ model provides for this, a service provider must be able to:

- Determine the path that IP routing takes for a particular customer's traffic
- Provision each router along the path for DiffServ
- Manually assure that not too many customers pass over that path, to avoid demand in excess of available bandwidth (the “over subscription” scenario)

Although this scenario is feasible with provisioning in a small network, a more scalable way to manage bandwidth is necessary to provide a point-to-point guarantee to the customer. Cisco DiffServ-Aware traffic engineering is ideal for this situation. By automatically choosing a routing path that satisfies the bandwidth constraint for each service class defined (such as premium, gold, silver, or bronze), DiffServ traffic engineering relieves the service provider from having to compute the appropriate path for each customer and each service class per customer.

Initial MPLS Network Topology

The hypothetical MPLS network scenario used in this solutions document is for a backbone representing a large network with a range of link speeds and congestions occurring at various links at different times. Figure 1 shows the logical connections between edge and core nodes in which guaranteed-bandwidth services are deployed. The provider-edge router terminates one side of the point-to-point connectivity to the customer edge. When connecting a new customer edge, configuration changes are required on both sides of the point-to-point connections between provider edge and customer edge, but no configuration changes are needed on the P nodes across the path.

Note: The solutions presented in this document are based on a hypothetical customer environment. All the IP addresses and configurations in this document are provided for illustrative purposes only.

Figure 1
Logical Connection of a Typical Service Provider Network

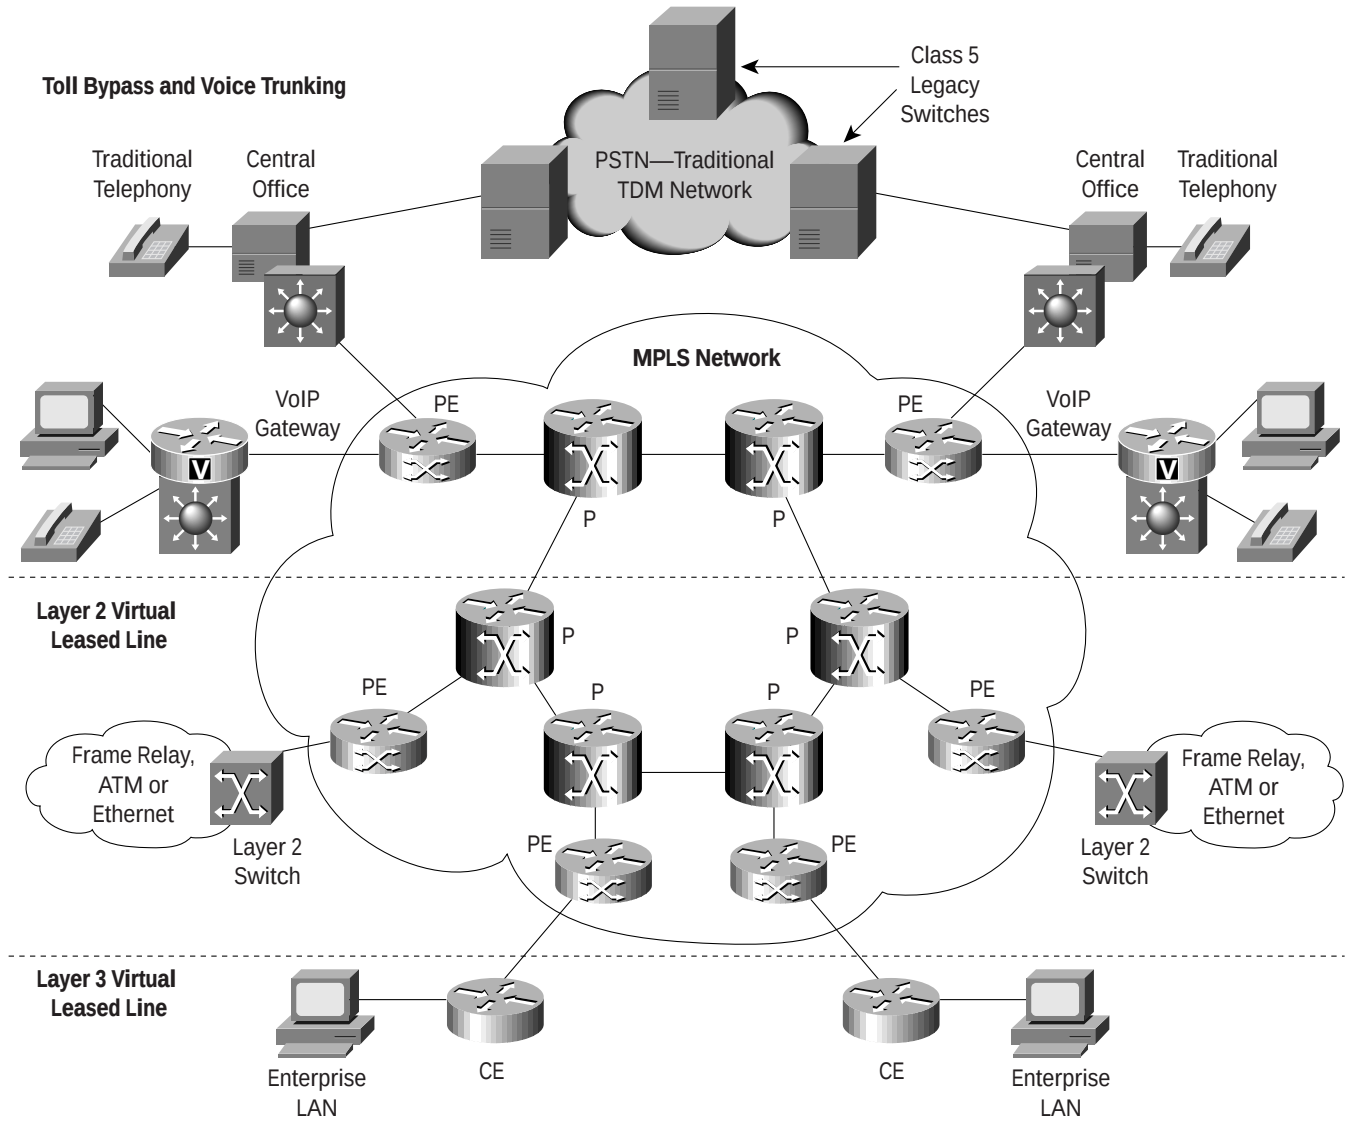

PE: Provider Edge
CE: Customer Edge
P: Provider Node

Proposed Solution: Guaranteed-Bandwidth Services with MPLS

Overall Guaranteed-Bandwidth Solution Overview

The strategy for implementing guaranteed-bandwidth services in a service provider network includes:

- Part 1: Overall guaranteed-bandwidth strategy
- Part 2: Implementation strategy for connecting end users into infrastructure

Note: To provide guaranteed-bandwidth services, the network can utilize the Cisco 7500 Series Router at the edge and the Cisco 12000 Series Internet Router at the edge and in the core of the network.

The features necessary to enable this solution exist in Cisco IOS Software starting with Release 12.0(14)ST.

Part 1: Overall Guaranteed-Bandwidth Strategy

Before the core network in Figure 1 can deploy guaranteed-bandwidth services for end users, MPLS must be implemented in the network. RSVP, originally described in RFC 2205, is the control and signaling protocol used in these guaranteed-bandwidth service scenarios.

Major Components

To deliver guaranteed-bandwidth services, a combination of features is required to support the overall solution. Features are categorized in two types: control-plane features and forwarding-plane features. Figure 2 below depicts these features. Provider-Edge Router: Cisco 12000 or 7500; Provider-Core Router: Cisco 12000.

Figure 2
Features in the Control Plane and Forwarding Plane

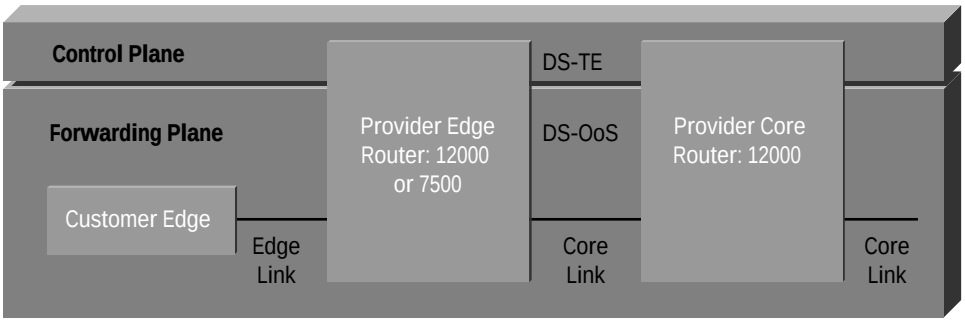

Features in the Control Plane

Cisco IOS Software delivers a powerful combination of industry-leading technology and features to build virtual leased-line and voice-trunking solutions with the assumptions and characteristics described above. The following Cisco IOS MPLS features are the essential ingredients in building a profitable and highly robust voice-trunking or toll-bypass trunking service.

Cisco MPLS traffic engineering—Cisco MPLS traffic engineering automatically sets up label-switched paths (LSPs) that can ensure, through appropriate QoS mechanisms, that the bandwidth, delay, and jitter constraints imposed by the toll-bypass or voice-trunking application are met. Additionally, MPLS traffic engineering is the first step to set up these paths for carrying voice traffic in a diverse manner for better network utilization, overall throughput, and resiliency in the network.

Cisco MPLS DiffServ-Aware Traffic Engineering—Traffic engineering in itself treats all traffic the same and does not differentiate among traffic types. In order to carry voice and data traffic on the same network, it may be necessary to account separately for the amount of voice traffic being transferred over the network, so as to provide the necessarily stricter QoS guarantees. Cisco DiffServ-Aware Traffic Engineering not only allows the configuration of a global pool for bandwidth accounting but also provides a restrictive subpool configuration for high-priority network traffic, such as voice. Both the available bandwidth on the global pool and the available bandwidth in the subpool are advertised by Interior Gateway Protocol (IGP) LSAs or time/length/values (TLVs), thus ensuring that each router

keeps track of the available bandwidth when admitting new LSPs for voice or high-priority traffic. In this manner, service providers, depending on their service-level agreement (SLA) requirements, can choose to overbook lower-priority classes or even underbook higher-priority traffic to meet tight QoS requirements. Obviously, they can charge a premium for that extra protection of voice traffic.

Cisco MPLS Fast Reroute (MPLS FRR)—Fast reroute is the ability to locally patch traffic onto a backup tunnel in case of a link or node failure with a fail-over time of 50 ms or less, a time that is competitive with Synchronous Optical Network automatic protection switching (SONET APS). Cisco FRR utilizes MPLS label stacking with RSVP signaling to create a backup tunnel around the link or node that needs to be protected. On detection of loss of signal from the link, the MPLS FRR application in Cisco IOS Software starts forwarding the traffic onto the backup tunnel, transparent to end users or applications such as VoIP or video, in 50 ms or less (actual failover time may be greater or less than 50ms, depending on the hardware platform, the number of TE Tunnels and/or Network prefixes).

Cisco MPLS AutoBandwidth Allocator—Cisco IOS Software supports another first: an MPLS traffic-engineering feature to ease constant monitoring and provisioning of the network; it is called Cisco AutoBandwidth Allocator. The AutoBandwidth feature constantly keeps track of average usage of the MPLS traffic-engineering tunnel and can resize traffic-engineering tunnel bandwidth to suit the traffic flow, thereby efficiently utilizing the available network bandwidth and maximizing profits for service providers. The average duration for which monitoring happens is configurable, thereby providing better control of network resources.

Features in the Forwarding Plane

Cisco IOS Software also provides a rich set of QoS features that are necessary to provide the minimum QoS guarantees to traffic-engineering tunnels. These mechanisms work hand-in-hand with DiffServ traffic engineering to provide a point-to-point guarantee for each service class.

To provide the QoS guarantees, the customer must control the following network characteristics:

- *Bandwidth guarantees*—Toll-bypass trunking requires the equivalent of an emulated circuit, point to point, in the network, with bandwidth guarantees. The network devices must be capable of scheduling traffic so that voice traffic always receives its share of the link capacity under any (moderate or heavy) congestion conditions. A full bandwidth guarantee for a service class is accomplished when the bandwidth reserved in the control plane, in the main pool, or in the subpool is equivalent to the bandwidth reserved in the forwarding plane for the same class, and the usage of the bandwidth for each reservation is compliant with the requested bandwidth.
- *Delay guarantees*—Bandwidth guarantees do not always ensure a proper delay or jitter guarantee. For example, a longer route across the network may provide a bandwidth guarantee but will not meet the delay requirement for tight QoS-based services. Applications such as voice trunking also require a delay guarantee. By carefully selecting the path across the network with traffic engineering, different classes can receive different minimum delay guarantees according to the service class.
- *Jitter bounds*—Voice-trunking applications require consistent and predictable network behavior. Network devices introduce jitter during traffic queuing and scheduling, regardless of how smooth the initial entry of traffic. Providing low network jitter also reduces the requirement of large de-jitter buffers in the end nodes, resulting in smooth playback of voice at the receiving end.
- *Advanced Cisco IOS QoS feature in the core and the edge, with DiffServ*—At the edge of the network, and before going into a tunnel, traffic is policed and colored appropriately. Coloring refers to marking the packets with the appropriate MPLS Type of Service bits value in the MPLS header. This color is then used in the core to identify the class to which the packet belongs. In the core, the Cisco Low-Latency Queuing (LLQ) scheme is deployed to

ensure the minimum bandwidth for tunnels of a particular class. This allows a service provider to ensure strict priority, and an assured amount of bandwidth for voice, while dividing the remaining bandwidth "pie" into slices (called Class-Based Weighted Fair Queuing, or CBWFQ) for the other tunnels and data traffic.

Configuration Tasks for Enabling Guaranteed-Bandwidth Services on the Core Network

To enable guaranteed-bandwidth services capability on the core network, perform the following general configuration tasks:

Step 1. Global Configuration

To enable global configuration, use the following global configuration commands:

- Enable traffic engineering on the router:

```
mpls traffic-eng tunnels
```

This command enables MPLS traffic engineering on a device. In order to use the feature, MPLS traffic engineering must also be enabled on the desired interfaces.

- Configure Open Shortest Path First (OSPF) or Integrated System-to-Integrated System (IS-IS) as the routing protocol:

Traffic Engineering requires a link-state routing protocol. Traffic-engineering supports OSPF or IS-IS for internal routing. The following is an example for IS-IS configuration; for more information on IP Routing Protocols Configuration go to: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/index.htm.

IS-IS example:

```
router isis
 redistribute static ip metric 0 metric-type internal level-2
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2
 net 49.0111.0000.0000.1001.00
 is-type level-2-only
 metric-style wide
 log-adjacency-changes
```

Note: Refer to the routing-protocol configuration manual for details on each command. Additional configuration on the interfaces may be required; for example, ip router isis is required as a minimum on the interface for IS-IS. Refer to routing-protocol configurations for details.

Note: MPLS forwarding requires ip cef to be enabled on the router.

Note: A loopback interface with /32 as the router ID is recommended for all control-plane configurations.

Step 2. Control-Plane Configuration on the Core Interfaces

To configure guaranteed-bandwidth services on the core interfaces, use the following two interface configuration commands:

- *Enable traffic engineering on the interfaces:*

```
mpls traffic-eng tunnels
```

This command enables the traffic-engineering tunnels on the interface. It configures the interface to send and receive RSVP signaling to establish traffic-engineering tunnels across this interface; both sides of the link need to have this configuration enabled.

- *Define the bandwidth allocation on the interfaces:*

```
ip rsvp bandwidth interface-kbps single-flow-kbps [sub-pool kbps]
```

This command enables RSVP reservations for traffic-engineering tunnels.

* `interface-kbps is` the amount of bandwidth (in kbps) on the interface that is available for reservation, and it is referred to as *global pool.*

* `single-flow-kbps` is the maximum amount of bandwidth (in kbps) allowed for a single flow. This parameter is ignored for traffic-engineering tunnel reservations.

* `[sub-pool kbps]` is the amount of bandwidth (in kbps) from the global pool available for reservations in a subpool.

Global pool and subpool indicate the bandwidth assigned for provisioning guaranteed-bandwidth services tunnels crossing the network. The bandwidth pools are provisioned on all the links inside the core network. When a new tunnel is admitted into the global-pool or the subpool bandwidth, the new tunnel cannot be built if the pool does not have enough bandwidth available to satisfy the new tunnel request.

The bandwidth allocated in this command represents the control-plane configuration for the traffic-engineering tunnels. The bandwidth assigned for global pool and subpool must be mapped with QoS configuration, as shown in the next step. To provide QoS guarantees for the bandwidth pools, each bandwidth pool should be assigned a QoS class and queuing parameters to provide the bandwidth guarantees on the core interfaces. The pool bandwidth must match the bandwidth assigned for the class on every link in the network in order to provide the bandwidth guarantees for the pool.

Note: The ability to define a subpool is a feature also referred to as Cisco DiffServ-Aware Traffic Engineering.

Step 3. Forwarding-Plane Configuration for the Core Interfaces

Forwarding-plane features for the core interfaces include congestion management and congestion avoidance.

Congestion management—With congestion management, packet classes are differentiated based on bandwidth and bounded delay. It is an automated scheduling system that provides bandwidth allocation to all network traffic based on its class. It also provides guaranteed minimum delay for certain classes carrying applications such as VoIP.

Note: The congestion management is achieved with Modified Deficit Round Robin (MDRR) on the Cisco 12000 and with CBWFQ on the Cisco 7500 (For more information regarding Configuring MDRR/WRED on a Cisco 12000 Series Router go to: [illegible] and for CBWFQ on the Cisco 7500 go to: [illegible])

Congestion avoidance—Congestion avoidance is achieved by Weighted Random Early Detection (WRED). Where packet classes are differentiated based on drop probability, WRED monitors network traffic, trying to anticipate and prevent congestion at common network and internetwork bottlenecks. WRED can selectively discard lower-priority traffic when an interface begins to get congested. It can also provide differentiated performance characteristics for different classes of service.

Step 4. Optional Core Configuration on the Core Interfaces

Fast reroute—This provides link protection for a link failure in less than 50-ms reroute time. This feature improves the availability of the traffic carried by traffic-engineering tunnels. For additional details on support and configuration of FRR on the core interfaces go to: [illegible]

Implementation Strategy for Connecting End Users into Infrastructure

Now that the core routers and links are configured for traffic engineering and QoS forwarding, perform the following steps to add end users and services:

Step 1. Create Traffic-Engineering Tunnel Interface

Customer-end traffic is transported as point-to-point traffic to the remote end users' site. End users' traffic is carried in traffic-engineering tunnels. Traffic-engineering tunnels provide the path selection and bandwidth reservation for the end users' traffic. Traffic-engineering tunnels are unidirectional, and hence two tunnels between the provider edge are required to establish two-way traffic flow.

The following sample configuration can be used to create a traffic-engineering tunnel between provider edges:

```
interface Tunnel0
 ip unnumbered Loopback0
 no ip directed-broadcast
 load-interval 30
 tunnel destination 11.11.14.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth sub-pool 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng fast-reroute
 tunnel mpls traffic-eng AutoBandwidth frequency 300 max-bw 4000 min-bw 1000
```

In the example, the Tunnel0 is being admitted in the subpool bandwidth. When a traffic-engineering tunnel is established, the network has admitted the path and bandwidth reservation into the network resources.

Note: As of Cisco IOS Software Release 12.0(16)ST, the current scalability of traffic-engineering tunnels is 600 tunnels at the headend router, 10,000 tunnels at the midpoint router, and 5000 tunnels at the tail-end router.

Tunnels can be created in different ways. Following is an analysis on implementation trade-offs:

1. One traffic-engineering tunnel on provider edge per customer edge and per point-to-point path—In this approach, new point-to-point connectivity between two customer edges requires a separate traffic-engineering tunnel provisioned with the bandwidth required.

- Advantages:
  - This configuration offers fine control and granularity on every connection admitted to the network.
  - The implementation procedure is simple: for every new connection, add a traffic-engineering tunnel and add the routing required (refer to Step 3).
  - At every node and every link, the network reflects the actual connections and bandwidth reservation for all the connections.
  - This approach is more suitable for midsize deployments with up to 100 tunnels per provider edge.
- Disadvantages:
  - This setup consumes provider-edge resources: every connection requires a traffic-engineering tunnel, consuming memory and interface description blocks (IDBs) on the routers.
  - Additional traffic-engineering tunnel aggregation methods are required to summarize the tunnels in the core.
  - A large number tunnels could result in a longer time for convergence and for network stabilization after rerouting.

2. Full-mesh traffic-engineering tunnel between provider edges—In this scenario, one tunnel per service class is established between each pair of provider-edge nodes, and this tunnel aggregates and carries all the traffic sent from the directly connected customer edges. Each tunnel is provisioned with enough bandwidth to support all the connections that cross it.

- Advantages:
  - A fewer number of tunnels results in less utilization of network resources.
  - Convergence and stabilization times are faster, even after a cold start.
- Disadvantages:
  - This approach offers less control over individual connections.
  - This approach requires estimating the amount of bandwidth to provision on each tunnel. Estimating bandwidth require ongoing adjustments to avoid over or under provisioning the tunnel bandwidth.

3. Full-mesh traffic-engineering tunnel between the P nodes—In this scenario, one tunnel per service class is established between Provider nodes, and this tunnel carries all the traffic sent from the downstream provider edges.

- Advantages:
  - Fewer tunnels in the core results in utilization of fewer network resources.
  - This approach provides aggregation in the core for tunnels established by the edge.
- Disadvantages:
  - This approach requires an additional layer of tunnels to carry guaranteed-bandwidth tunnels.
  - This approach involves the added complexity of managing additional tunnels.

4. Traffic-engineering tunnel headend located on the customer-edge router—In this scenario, the tunnel headends are established at the customer edges. The customer edge then signals the tunnel across the provider edge and the network.

- Advantages:
  - Resources are freed from the provider edge by offloading the tunnel headend to the customer edge.
  - Provider-Edge scalability now is the same as that for the P node, up to the 10,000-tunnel midpoint.
  - New connections do not require additional configuration on the provider edge; the tunnel headend is added on the customer edge.
  - The customer edge can create on-demand guaranteed-bandwidth requests.
  - Bandwidth can be provisioned to extend to and include the edge link.
- Disadvantages:
  - This solution makes the most sense in a managed-customer edge scenario.
  - No specific authentication mechanism is possible in RSVP.
  - Billing mechanisms still need to be identified.
  - OSPF and IS-IS are scalable, because now they have to run between the customer edge and the provider edge.

Additional Tunnel Configuration Options to Consider

Tunnel priority—The command tunnel mpls traffic-eng priority establishes a priority level for this tunnel, so that if lower-priority tunnels are already established on the same path, the lower-priority tunnel is forced to reroute into another path, and clear the path for the higher-priority tunnel. Setting a higher priority for the tunnel carrying traffic such as VoIP can provide precedence for VoIP over data traffic, for example, and enable the VoIP traffic to use a shorter path; for example, a path with a lower hop count. The result could be lower end-to-end delay.

AutoBandwidth—If the AutoBandwidth option is used, the bandwidth reserved for the tunnel is automatically adjusted within the same main or subpool. The adjusted bandwidth value reflects the peek traffic rate passing through the tunnel in the previous sampling period. The new bandwidth value is bound within the range of a maximum or minimum bandwidth configuration. When traffic going through the tunnel is low, the tunnel adjusts its configuration in the next period to lower bandwidth reservation and release bandwidth resources back to the network. In the same way, when traffic going though the tunnel increases, the tunnel requests additional bandwidth resources from the network at the next period up to the configured maximum bandwidth. For more information regarding DiffServ-Aware Traffic Engineering go to: [illegible]

Fast reroute—Using FRR, the tunnel is protected when link protection is used anywhere on the path. Link protection helps improve the network availability when a failure occurs on a protected link. For additional information for the FRR documentation go to: [illegible]

Path option—Path option is the method by which the tunnel determines which path to use across the network. The previous example uses dynamic path routing. With the option dynamic, the provider edge uses the IGP link-state database to determine the best path with available bandwidth for reservation. The path-option method requires an additional step of specifying an explicit list of IP addresses for the tunnel to cross.

Step 2. Configure Interface to the End User

The configuration of the interface to the end user will vary, depending on the media and encapsulation. The common configuration required, however, includes packet classification, policing, and shaping.

Packet classification and marking—Packets need to be classified at the edge of the network before labels are assigned. Packets are classified according to the service class they belong to. Although the classification is often configured on the inbound interface of the provider edge, it is recommended that the classification function be performed on the customer edge in managed-customer edge environments, in order to offload additional work from the provider edge.

Note: The packet classification is achieved with committed access rate (CAR) on the Cisco 12000, and with a modular QoS command-line interface (CLI), it uses the set and police commands within the modular QoS command line (MQC) on the Cisco 7500. Traffic can also be marked using the MQC on Supervisor Engine 3 line cards on the Cisco 12000. For more information for Configuring MDRR/WRED on a Cisco 12000 Series Router go to: [redacted] and for RSVP-Traffic Engineering Extensions to RSVP for LSP Tunnels go to: draft-ietf-mpls-rsvp-lsp-tunnel-08.txt:.

Packet policing and shaping—Cisco IOS QoS offers two kinds of traffic regulation mechanisms—policing and shaping. These features are used on the inbound traffic to ensure that packets adhere to the requested bandwidth. The policing function checks to see whether incoming traffic conforms to or exceeds the requested bandwidth, and then re-marks or drops the packet accordingly. The shaping function typically delays excess traffic using a buffer, or a queuing mechanism, to hold and shape the flow when the data rate of the source is higher than expected. Traffic policing and shaping can work in tandem. Traffic policing is required to enforce the requested bandwidth utilization for each point-to-point connection. Refer to the section "Implementation of Proposed Solution: Guaranteed Bandwidth with MPLS" for a configuration example. This function should be configured on the inbound interface of the provider edge, in a managed-customer edge scenario. It should be performed on the customer edge in order to offload additional work from the provider edge.

Note: Packet policing and shaping are achieved with CAR on the Cisco 12000, and with modular QoS CLI, it uses the `set` and `police` commands within the MQC on the Cisco 7500. Traffic can also be marked using the MQC on Supervisor Engine 3 line cards on the Cisco 12000. For more information for Configuring MDRR/WRED on a Cisco 12000 Series Router go to: [redacted] and for RSVP-Traffic Engineering Extensions to RSVP for LSP Tunnels go to: draft-ietf-mpls-rsvp-lsp-tunnel-08.txt:

Packet queuing—Cisco routers running Cisco IOS Software have numerous different queuing mechanisms. Queuing on routers is necessary to accommodate bursts when the arrival rate of packets is greater than the departure rate. Queuing is used to ensure minimum bandwidth guarantees for each class and to provide low-latency queuing for VoIP traffic.

Step 3. Forwarding Traffic onto Tunnel

At this step, traffic received from the customer-edge router needs to be forwarded on the assigned traffic-engineered tunnel. Traffic can be forwarded onto the traffic-engineering tunnel in two ways: static routing and policy-based routing (PBR).

Static routing—Static routing is the recommended and the simplest method for forwarding traffic onto the tunnel. The forwarding is based on the destination prefix address. Multiple static routes can be used for multiple prefixes to the same destination.

Policy-based routing—PBR is the most flexible method for forwarding traffic onto the tunnel. The forwarding decision is based on multiple criteria existing in the extended access control list (ACL): destination prefix, source prefix, class of service, port number, and source interface. PBR is used to forward traffic when static routing is not sufficient, for example, when there is need to specify a class or a need to be specific on the source interface allowed to use the tunnel. For more information on policy-based routing go to: [illegible]

Note: At the time this document was created, PBR is not supported for virtual-private-network (VPN) traffic routed from a virtual routing and forwarding (VRF) interface.

Overall Guaranteed-Bandwidth Prerequisites

Before implementing guaranteed-bandwidth services, the network must meet the following prerequisites:

- The network must have an IP address-allocation plan.
- The network must have a loopback interface with mask /32 to identify each provider-edge router.
- The network must use a link-state protocol (OSPF or IS-IS) in the core as an IGP.
- The network must have an allocation plan for multiple classes of service for DiffServ QoS.
- The network must support software and hardware for Cisco DiffServ-Aware Traffic Engineering.
- The network must have enough to memory the support the network scaling.

Overall Guaranteed-Bandwidth Network Topology

Figure 3 shows the traffic-engineering tunnel established between two provider edge routers to provide guaranteed bandwidth connectivity.

Figure 3
End-To-End Guaranteed-Bandwidth Tunnels

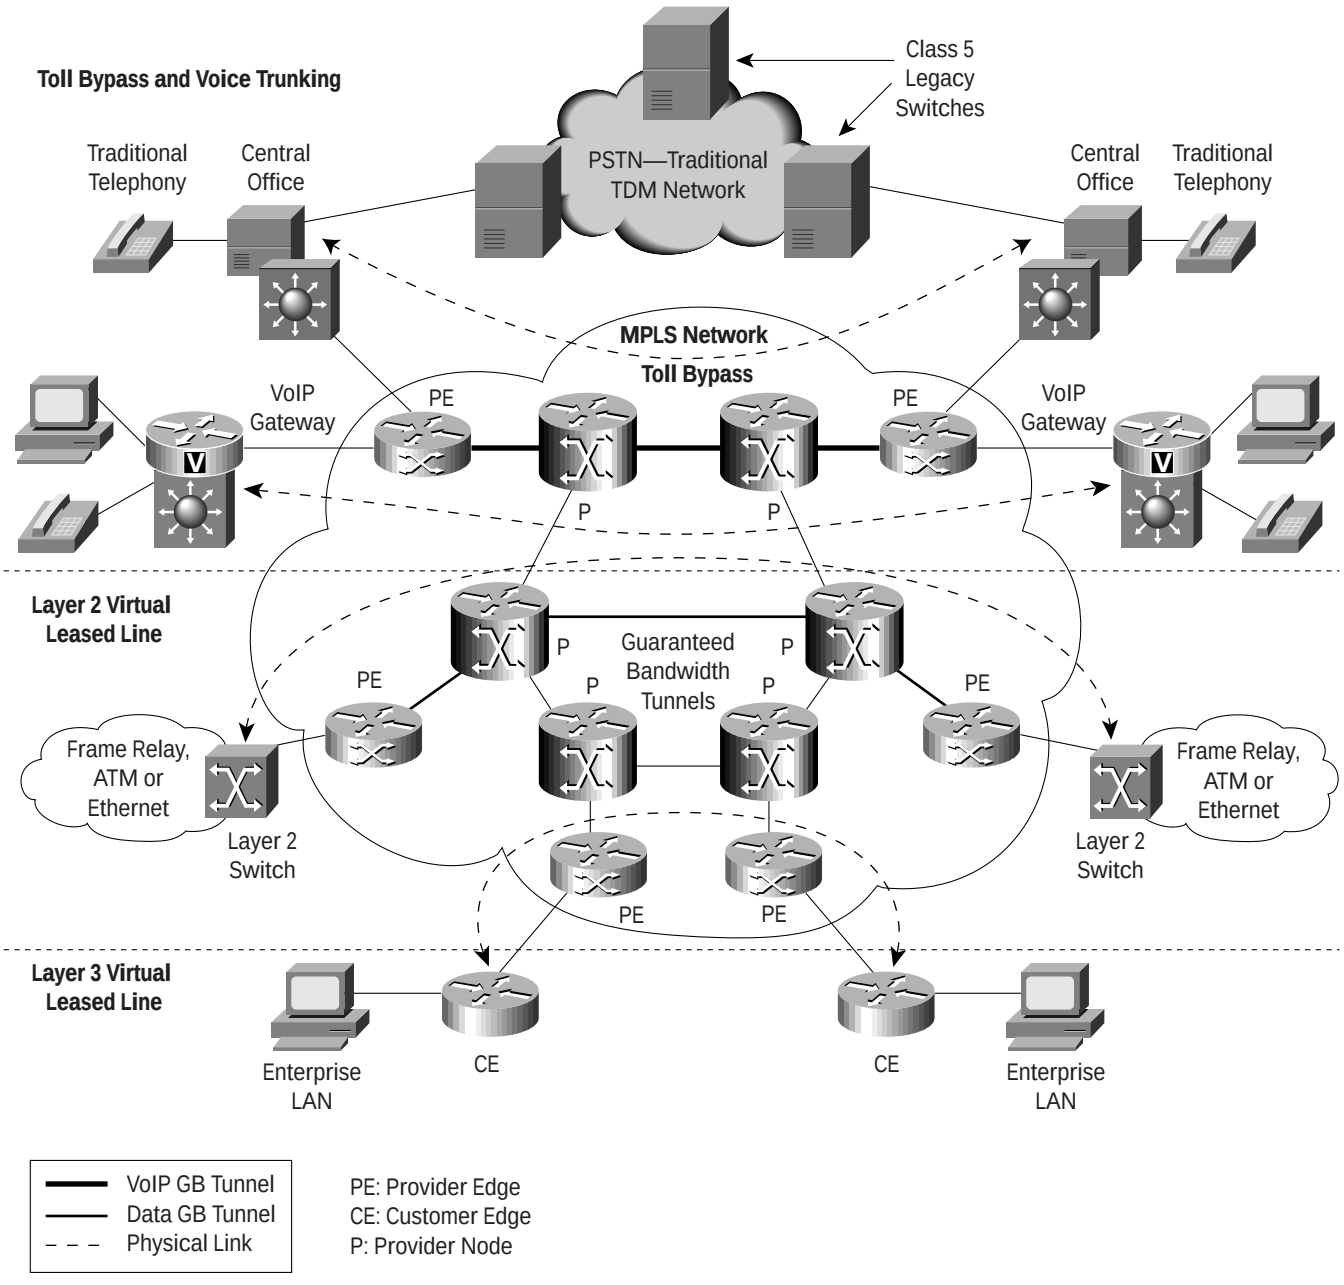

Overall Guaranteed-Bandwidth Services Benefits

The benefits of using traffic engineering with QoS to provide guaranteed-bandwidth services are as follows:

- By deploying the guaranteed-bandwidth services solution, customers now can offer voice and data services with point-to-point bandwidth guarantees across multiple link types and different link speeds using the MPLS network.
- This feature offers new premium services for high-priority traffic, such as voice traffic or online transaction processing with tight guarantees for throughput, delay, and more, and best-effort traffic on the same network.
- Achieve higher network availability by using MPLS FRR to quickly use alternate traffic-engineered paths—in 50 ms or less. And achieve higher network resource efficiency by forwarding packets on multiple traffic-engineering paths.

Overall Guaranteed-Bandwidth Ramifications

The ramifications for using MPLS and traffic engineering to implement guaranteed-bandwidth services are as follows:

- The traffic-engineering tunnel state must be maintained in the router. This situation uses additional memory resources in the edge router.
- Routers that utilize multiple QoS features may experience additional load on CPU resources.

Implementation of Proposed Solution: Guaranteed-Bandwidth Services with MPLS

This section describes implementing guaranteed-bandwidth services and connecting customers into the infrastructure of an IP backbone. It covers the following scenarios:

- Implementing the core configuration
- IP virtual leased-line scenario
- Layer 2 virtual leased-line scenario
- Voice-trunking and toll-bypass scenario

Implementing the Core Configuration for Guaranteed-Bandwidth Services

- This section describes the implementation of the core configuration necessary to enable guaranteed-bandwidth services.

Guaranteed-Bandwidth Services Core Configurations Strategy

For the core network to carry guaranteed-bandwidth services, the customer needs to partition the network resources in the control and forwarding planes. The control-plane resources are partitioned by specifying the network bandwidth capacity allocated for different traffic-engineering bandwidth pools. The forwarding-plan resources are partitioned by specifying QoS forwarding characteristics on all the core links to handle multiple classes of service and characteristics such as minimum bandwidth guarantees, packet discards, and queuing priority.

Guaranteed-Bandwidth Services Core Configuration Topology

Figure 4 shows a diagram of a network capable of delivering guaranteed-bandwidth services.

Figure 4
Network Diagram for Guaranteed-Bandwidth Core

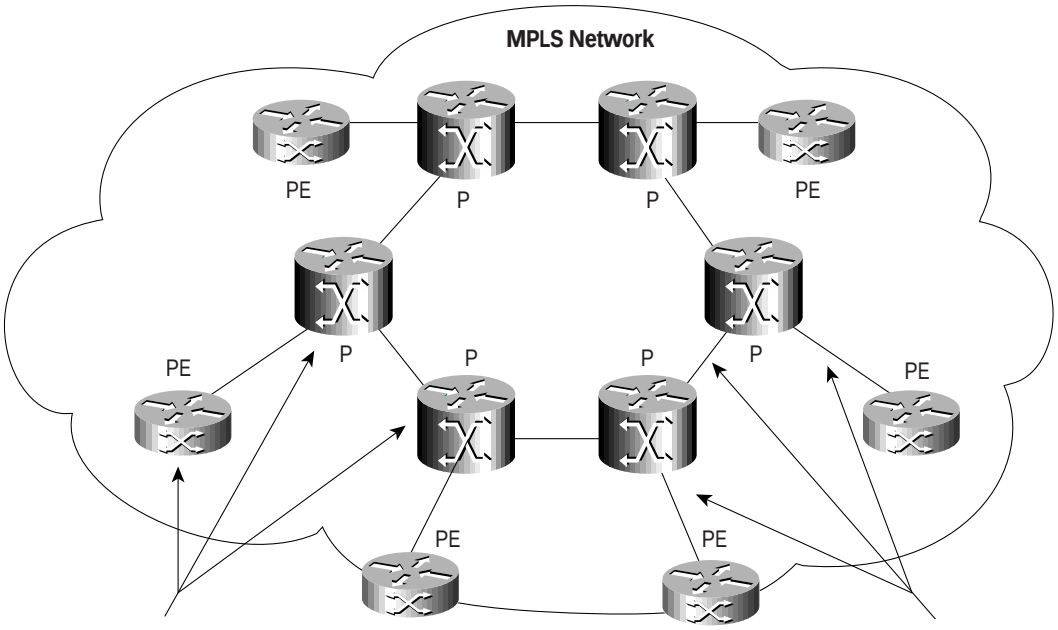

DiffServ Quality of Service and DiffServ-Aware Traffic Engineering

PE: Provider Edge
CE: Customer Edge
P: Provider Node

Guaranteed Bandwidth Services Core Configurations Benefits

The benefits of using guaranteed bandwidth services are as follows:

- This solution allows multiple traffic classes to run on the network without having to operate separate networks.
- Customers can reserve bandwidth, determine the bandwidth reserved in real time, and provision new services on line.

Guaranteed-Bandwidth Services Core Configuration Ramifications

Customers who partition the network resources must perform additional configuration.

Guaranteed-Bandwidth Services Core Configuration Summary

The following is a summary of the tasks to configure the core network for guaranteed-bandwidth services. In this example, the customer has two traffic-engineering bandwidth pools. The bandwidth pools have two class-of-service queues reserved, Class 4 and Class 5.

Step 1. Configure guaranteed-bandwidth services bandwidth pools.

In the following example, the core link capacity has 100 Mbps allocated for the main pool, and 20 Mbps of the main pool is allocated for the subpool.

```
Interface POS1/0
description OC3 link
mpls traffic-eng tunnels
ip rsvp bandwidth 100000 10000 sub-pool 20000
```

Step 2. Configure class of service for guaranteed-bandwidth services.

On the Cisco 12000, to allocate bandwidth on the links for the guaranteed-bandwidth services pools, configure MDRR on the interfaces by creating an MDRR cos-queue-group, as shown in the following example:

```
cos-queue-group oc3-link
precedence 4 queue 4
precedence 5 queue low-latency
queue 4 500 !(Queue 4 has weight value of 500)
queue low-latency strict-priority !(low-latency queue is in strict-priority mode)
```

To specify the packet-drop characteristics for the guaranteed-bandwidth services classes, configure the WRED parameters as shown in the following example:

```
cos-queue-group oc3-link
random-detect-label 4 500 1250 1
precedence 4 random-detect-label 4
exit
```

To map the cos-queue-group on the interface, use the following commands:

```
interface POS1/0
tx-cos oc3-link
```

This configuration defines the bandwidth guarantees and packet-discard characteristics for Classes 4 and 5 on the OC-3 link. Class 5 is assigned low-latency queuing in order to receive guaranteed delay for voice-grade traffic. This example applies to the Cisco 12000 with pre-MQC support. For details on configuring congestion management and congestion avoidance with MQC go to

Note: The parameters used in this example are for illustration only. The actual bandwidth reserved and packet discards for each class are dependent on the other class configurations. For details on configuring MDRR and WRED go to:

Note: Class 5 with low-latency queuing does not have early discard configured with WRED.

Step 3. Verify the MDRR/WRED statistics and configurations:

```
ios-gsr8b#show interfaces pos 1/0 random-detect
POS1/0
cos-queue-group: oc3-link
RED Drop Counts
 Tx Link To Fabric
RED Label Random Threshold Random Threshold
0 0 0 0 0
1 0 0 0 0
2 0 0 0 0
3 0 0 0 0
4 72327 322476 0 0
```

```
5 0 0 0 0
6 0 0 0 0
Tx-queue-limit drops: 0

Queue Lengths
Tx Queue (DRR configured) oc3-link
Queue Average High Water Mark Weight
0 0.000 0.000 10
1 0.000 0.000 10
2 0.000 0.000 10
3 0.000 0.000 10
4 846.000 956.000 500
5 0.000 0.000 10
6 0.000 0.000 10
Low latency 128.000 214.000 100

Tx RED config
 Precedence 0: not configured for drop
 Precedence 1: not configured for drop
 Precedence 2: not configured for drop
 Precedence 3: not configured for drop
 Precedence 4: 500 min threshold, 1250 max threshold, 1/1 mark weight
 Precedence 5: 500 min threshold, 1250 max threshold, 1/1 mark weight
 Precedence 6: not configured for drop
 Precedence 7: not configured for drop
weight 1/2
```

Verify IP RSVP configurations:

```
ios-gsr8b#show ip rsvp interface
interface allocated i/f max flow max sub max
PO1/0 0G 420M 50M 20M
Fa6/0 25M 50M 35M 20M
```

Implementing IP Virtual Leased-Line Scenario

This section describes the implementation of IP virtual leased-line (VLL) service connectivity.

IP Virtual Leased-Line Definition

The primary purpose of this service is to transport IP in a point-to-point manner. Connectivity between the edge device and provider-edge router is, therefore, always an IP connection. This IP trunk may emulate a voice trunk or simply transport data between a backup site and a data center. In each case, the QoS requirements are distinct. A strict VLL implementation provides not only connectivity, but also point-to-point guarantees of bandwidth, jitter, delay, packet drops, and improved availability.

IP Virtual Leased-Line Strategy

In this scenario, a bandwidth pool is allocated on all the core links for the VLL service. The VLL pool could be the main pool or the subpool in the traffic engineering. Bandwidth guarantees are provided in the network by admitting additional services only if bandwidth is available in the VLL bandwidth pool.

In a simple example—point-to-point connectivity between two sites—the bandwidth admitted for this connection into the VLL bandwidth pool is equal to the access rate of the customer-edge site, guaranteeing the reservation of full bandwidth between the two sites.

The bandwidth reserved for a connection, however, can be rate limited to a fraction of the access speed, using QoS features and techniques such as MQC policer, CAR for backward compatibility, and traffic shaping. Only the required bandwidth is reserved for the service.

When the point-to-point connection is rate limited to a fraction of the access speed, additional VLL connections can be added to the access interface. In order to provide guaranteed bandwidth, however, the total bandwidth reserved for all the connections should not exceed the customer-edge access speed.

IP Virtual Leased-Line Network Topology

Figure 5 shows an IP VLL connection with a guaranteed-bandwidth tunnel over an MPLS backbone.

Figure 5
Network Diagram for IP Virtual Leased Line

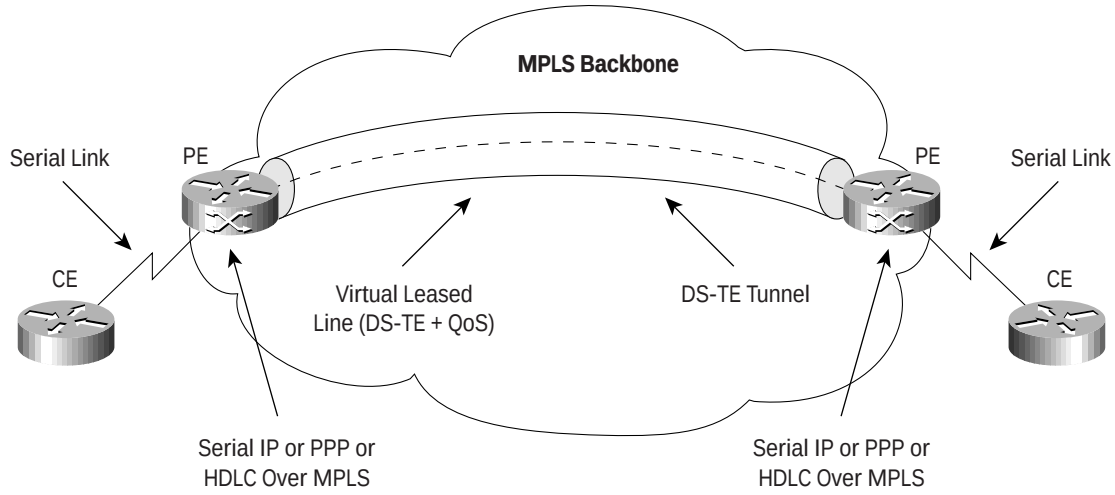

IP Virtual Leased-Line Benefits

The benefits of deploying Layer 2 VLL are as follows:

- Layer 2 VLL offers new premium services for high-priority traffic with tight QoS guarantees for throughput and delay.
- Higher network availability to Layer 2 circuits is achieved by using Cisco MPLS FRR to quickly use alternate traffic-engineered paths—in 50 ms or less.
- Layer 2 VLL provides point-to-point bandwidth admission and guarantees for the IP traffic.
- Layer 2 VLL reduces costs by taking advantage of available bandwidth on alternative paths; it also provides guarantees for high-priority traffic.
- Layer 2 VLL offers the ability to provision IP VLL service capacity on the network and manage core links utilization.

IP Virtual Leased-Line Ramifications

The ramifications of deploying the IP VLL scenario are as follows:

- It is more complex to implement and troubleshoot than an IP network.

- Extensive provisioning and automation tools are necessary to enable large-scale deployment.

IP Virtual Leased-Line Configuration Summary

The following is a summary of tasks to perform to configure a provider-edge router for IP VLL.

Step 1. Admit the VLL onto the traffic-engineering bandwidth pool.

The following configuration reserves 10 Mbps into the subpool bandwidth.

```
interface tunnel 10
 tunnel mpls traffic-eng bandwidth sub-pool 10000
```

In this step, the traffic-engineering tunnel headend is configured with a traffic-engineering tunnel for the IP VLL connection and is using the subpool. Refer to the "Create Traffic Engineering Tunnel Interface" step for a discussion on traffic-engineering headend location alternatives.

Step 2. Configure for policing and marking of the IP VLL traffic at the inbound provider edge.

```
interface Ethernet5/0
 rate-limit input 10000000 10000000 10000000 conform-action set-prec-transmit 5
exceed-action drop
```

This configuration marks the incoming traffic with IP Precedence to the QoS class assigned for the IP VLL on the network, and polices the traffic rate to 10 Mbps by dropping any excess traffic. This example applies to the Cisco 12000 with pre-MQC support. For details on configuring policing and marking with MQC go to http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt8/qcfmdcli.htm.

Step 3. Forward the traffic onto the tunnel.

The following configuration at the tunnel headend router forwards the IP VLL traffic onto the provisioned traffic-engineering tunnel using static route.

```
ip route 172.120.35.0 255.255.255.0 tunnel10
```

In this example, the static route command forwards the IP VLL traffic onto the traffic-engineering tunnel. In cases where the headend router will be sending other IP traffic to the same destinations that are not part of the IP VLL connection, IP PBR with extended ACLs can be used to distinguish between traffic sources, and forward the traffic appropriately.

Note: After the traffic has entered the ingress provider edge and has been marked, the outbound physical interface on the provider edge and every hop along the path will treat the traffic based on its marking and the interface configuration shown in Step 2 in the section "Guaranteed-Bandwidth Services Core Configuration Summary."

Step 4. Verify the configuration.

Verify that the traffic is forwarded onto the tunnel.

```
ios-gsr8b#show interface tunnel 0
Tunnel10 is up, line protocol is up
 Hardware is Tunnel
 Interface is unnumbered. Using address of Loopback0 (11.11.12.1)
 MTU 1514 bytes, bandwidth 9 Kbit, DLY 500000 usec, rely 255/255, load 138/255
 Encapsulation TUNNEL, loopback not set
 Keepalive set (10 sec)
 Tunnel source 11.11.12.1, destination 11.11.14.1
 Tunnel protocol/transport Label Switching, key disabled, sequencing disabled
 Checksumming of packets disabled
 Last input never, output 7w0d, output hang never
 Last clearing of "show interface" counters 6w3d
```

```
 Queueing strategy: fifo
 Output queue 0/0, 0 drops; input queue 0/75, 0 drops
 30 second input rate 0 bits/sec, 0 packets/sec
 30 second output rate 294000 bits/sec, 26 packets/sec
 0 packets input, 0 bytes, 0 no buffer
 Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
 1025 packets output, 1420950 bytes, 0 underruns
 0 output errors, 0 collisions, 0 interface resets
 0 output buffer failures, 0 output buffers swapped out
```

Verify that the tunnel is established:

```
ios-gsr8b#sh mpls traffic-eng tunnels tunnel 10
Name: ios-gsr8b_t1 (Tunnel10) Destination: 11.11.14.1
 Status:
 Admin: up Oper: up Path: valid Signalling: connected
 path option 5, type dynamic (Basis for Setup, path weight 20)
 Config Parameters:
 Bandwidth: 10000 kbps (Sub) Priority: 7 7 Affinity: 0x0/0xFFFF
 AutoRoute: disabled LockDown: disabled Loadshare: 10000 bw-based
 AutoBandwidth: disabled(0/1) 0 Bandwidth Requested: 10000
 InLabel : -
 Outzabel : FastEthernet6/0, 18
 RSVP Signalling Info:
 Src 11.11.12.1, Dst 11.11.14.1, Tun_Id 1, Tun_Instance 939
 RSVP Path Info:
 My Address: 11.11.2.1
 Explicit Route: 11.11.2.2 11.11.3.2 11.11.3.1 11.11.14.1
 Record Route: NONE
 Tspec: ave rate=10000 kbits, burst=1000 bytes, peak rate=10000 kbits
 RSVP Resv Info:
 Record Route: NONE
 Fspec: ave rate=10000 kbits, burst=1000 bytes, peak rate=10000 kbits
 History:
 Tunnel:
 Time since created: 49 days, 3 hours, 19 minutes
 Time since path change: 10 days, 3 hours, 6 minutes
 Current LSP:
 Uptime: 10 days, 3 hours, 6 minutes
 Selection: reoptimation
 Prior LSP:
 ID: path option 5 [76]
 Removal Trigger: path verification failed
```

Verify the label imposition and forwarding:

```
ios-gsr8b#show ip cef 10.10.10.1
10.10.10.0/24, version 267, attached
0 packets, 0 bytes
 tag information set
 local tag: tunnel head
 fast tag rewrite with Tu1, point2point, tags imposed {18}
 via Tunnel1, 0 dependencies
 valid adjacency
 tag rewrite with Tu1, point2point, tags imposed {18}
```

Implementing Layer 2 Virtual Leased-Line Scenario

This section describes the implementation of Layer 2 VLL service connectivity.

Layer 2 Virtual Leased-Line Definition

This service focuses on transporting Layer 2 protocols such as Ethernet, Frame Relay, and ATM in a point-to-point fashion across an MPLS networks. Layer 2 transport across an MPLS network may be required either to extend existing services or to provide simple, easy-to-provision services that are attractive to enterprise customers. For example, one service gaining popularity with providers is Ethernet over MPLS.

With Layer 2 VLL, customers can trunk non-IP protocols such as AppleTalk and Internetwork Packet Exchange (IPX) across the provider cloud, or extend virtual-LAN (VLAN) domains by transporting raw Ethernet frames. Service providers can use this service to create remote peering points that appear as a single hub by extending the broadcast domains and trunking Ethernet. Another example is to provide services to multidwelling units by providing Ethernet connect and then trunking the Ethernet to the point of presence (POP) without adding any routing or content services at the customer location.

Layer 2 Virtual Leased-Line Strategy

In this scenario, the customer transports the Layer 2 circuit onto a traffic-engineering tunnel. To transport the Layer 2 circuit, the ingress provider edge maps the destination IP address of the egress endpoint of the circuit to a traffic-engineering tunnel with a provisioned bandwidth. The bandwidth of the Layer 2 circuit is admitted into the main pool or the subpool, depending on the service class.

The traffic-engineering tunnel headend could be established on the router encapsulating the Layer 2 frames with traffic-engineering support, or on a separate and dedicated upstream router. The following Layer 2 VLL scenario is implemented by transporting Any Transport over MPLS (AToM). The network is configured with one Class of Service reserved for the VLL service.

Layer 2 Virtual Leased-Line Network Topology

Figure 6 shows a network diagram of Layer 2 VLL with AToM.

Figure 6
Network Diagram for Layer 2 VLL

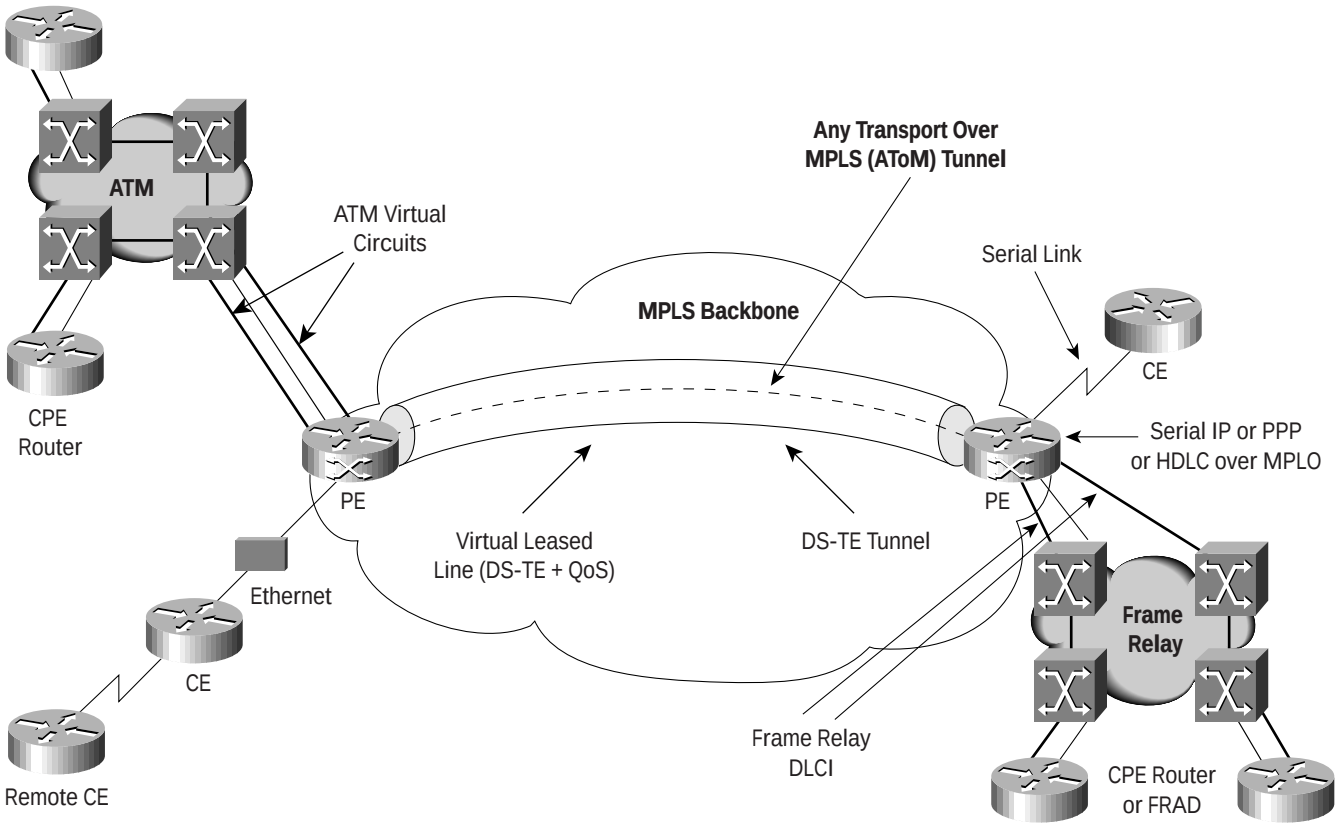

Layer 2 Virtual Leased-Line Benefits

The benefits of deploying Layer 2 VLL are as follows:

- Layer 2 VLL uses the MPLS network to provide Layer 2 connectivity with QoS guarantees.
- Layer 2 VLL achieves higher network availability to Layer 2 circuits by using Cisco MPLS FRR to quickly use alternate traffic-engineered paths—in 50 ms or less.
- Layer 2 VLL provides point-to-point bandwidth admission and guarantees for the Layer 2 traffic.
- Layer 2 VLL reduces costs by taking advantage of available bandwidth; it also provides guarantees for high-priority traffic.
- Layer 2 VLL offers the ability to provision Layer 2 service capacity on the network and manage core links utilization.

Layer 2 Virtual Leased-Line Ramifications

The ramifications of deploying the Layer 2 VLL scenario are as follows:

- A Layer 2 VLL network is more complex to implement and troubleshoot than an IP network.
- Extensive provisioning and automation tools are necessary to enable large-scale deployment.

Layer 2 Virtual Leased-Line Configuration Summary

The following is a summary of tasks to configure the provider-edge router for Layer 2 transport with Ethernet over MPLS. For other AToM media types, similar concepts apply.

Step 1. Create a Layer 2 circuit on the ingress provider edge.

```
!
mpls label protocol ldp
!
interface loopback0
 ip add 10.1.1.2 255.255.255.255
!
interface Vlan25
 no ip address
  mpls l2transport route 10.1.1.3 2000
!
```

The value “2000” is the virtual-circuit ID. It must be unique for each virtual circuit. The virtual-circuit ID is used to connect the endpoints of the connection.

Note: In the case of one connection per tunnel, a separate loopback needs to be created and associated in the Layer 2 route-map statement. The same loopback can be used for all connections if there are multiple connections on the same tunnel between two Layer 2 endpoint peers.

Step 2. Configure QoS for the Layer 2 circuit.

The following configuration utilizes the MQC to mark the traffic received from the VLAN interface.

```
!
class-map blue
 match any
!
policy-map badger
 class blue
 set mpls experimental 5
 shape average 2000000 8000 8000
!
interface vlan25
 no ip address
 service-policy input badger
!
```

This configuration marks the traffic received from Vlan25 into Class 5, the subpool class. Traffic received on this interface is also shaped into an average 2 Mbps.

Note: This example utilizes input traffic shaping with MQC supported on the Cisco 7600 to accomplish policing. For policing on other platforms, the MQC policer should be used.

Step 3. Admit the Layer 2 circuit bandwidth to the traffic-engineering pool.

The following configuration reserves the 2 Mbps into the subpool, matching the traffic shaper on the VLAN interface.

```
Interface tunnel 10
 Tunnel mpls traffic-eng bandwidth sub-pool 2000
```

In this example, the traffic-engineering tunnel headend is configured with a tunnel for every connection, where a new tunnel is created for every new connection and the bandwidth on the tunnel is set to reflect the actual connection reservation.

In order to maintain a separate traffic-engineering configuration for each Layer 2 connection, each Layer 2 connection must have a different Layer 2 egress endpoint destination IP address, even if multiple connections are targeted to the same Layer 2 egress endpoint. The destination IP addresses are created by loopback interfaces on the egress endpoint. This approach may encounter scaling limitations because of the required additional loopback interfaces and because a new tunnel interface is required for every new connection.

An alternative approach would be to carry multiple Layer 2 connections on a dedicated tunnel to the same egress provider edge with same service class. The dedicated tunnel bandwidth is provisioned with the additional capacity to support adding the new tunnels. Also, multiple connections per tunnel utilize the same destination IP address for the Layer 2 egress endpoint, making it more scalable and easier to manage.

Note: Locating the tunnel headend on the Layer 2 ingress provider edge requires traffic-engineering and DiffServ traffic-engineering support. If the Layer 2 ingress provider edge does not support traffic engineering, an upstream router with traffic-engineering support can provide the VLL guarantees similarly.

Step 4. Forward the Layer 2 traffic onto the traffic-engineering tunnel.

The following configuration forwards the Layer 2 traffic onto the provisioned traffic-engineering tunnel using static route.

```
ip route 10.1.1.3 255.255.255.255 tunnel10
```

The destination address used in the static route is the Layer 2 egress endpoint. This command is configured on the traffic-engineering headend router.

Note: After the traffic has entered the ingress provider edge and has being marked, the outbound physical interface on the provider edge and every hop along the path will treat the traffic based on its marking and the interface configuration shown in Step 2 in the section “Guaranteed-Bandwidth Services Core Configuration Summary.”

Step 5. Verify the configuration.

To make sure the label forwarding table is built correctly, check the MPLS forwarding table:

```
Router# show mpls forwarding-table
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
16 Untagged 10.255.254.254/32 0 Vl2 192.168.0.1
17 Pop tag 172.30.0.0/16 0 Gi6/3 172.16.0.1
18 Pop tag 172.20.0.0/16 0 Gi6/3 172.16.0.1
19 148 172.29.0.0/16 0 Gi6/3 172.16.0.1
20 77 172.20.0.1/32 6308338115 Gi6/3 172.16.0.1
23 Untagged EoMPLS(4) 94538 Vl4 point2point
24 Untagged EoMPLS(101) 847 Vl101 point2point
```

To view the state of the currently routed virtual circuits, check the connections table:

```
Router# show mpls l2transport vc
Transport Client VC Local Remote Tunnel
VC ID Intf State VC Label VC Label Label
4 Vl4 UP 23 21 77
101 Vl101 UP 24 22 77
```

Add the keyword detail to see detailed information about each virtual circuit.

```
Router# show mpls l2transport vc detail
VC ID: 4, Local Group ID: 25, Remote Group ID: 17 (VC is up)
Client Intf: Vl4 is up, Destination: 172.21.0.1, Peer LDP Ident: 172.20.0.1:0
Local VC Label: 23, Remote VC Label: 21, Tunnel Label: 77
Outgoing Interface: Gi6/3, Next Hop: 153.1.0.1
Local MTU: 1500, Remote MTU: 1500
Imposition: LC Programmed
Current Imposition/Last Disposition Slot: 6/32
Packet Totals(in/out): 1334/1337
Byte Totals(in/out): 95248/100812

VC ID: 101, Local Group ID: 27, Remote Group ID: 19 (VC is up)
Client Intf: Vl101 is up, Destination: 172.21.0.1, Peer LDP Ident: 172.20.0.1:0
Local VC Label: 24, Remote VC Label: 22, Tunnel Label: 77
Outgoing Interface: Gi6/3, Next Hop: 153.1.0.1
Local MTU: 1500, Remote MTU: 1500
Imposition: LC Programmed
Current Imposition/Last Disposition Slot: 6/32
Packet Totals(in/out): 11/6211757
Byte Totals(in/out): 847/2065861499
```

Implementing Voice-Trunking and Toll-Bypass Scenario

- This section describes the implementation of voice-trunking and toll-bypass services.

Voice-Trunking and Toll-Bypass Definition

This service focuses on integrating data and voice services on the same packet-switched network. Toll bypass is the ability to provide hard bandwidth, delay, jitter, and loss guarantees for voice traffic carried from a central office or packet-based telephone switch onto an IP infrastructure. Voice trunking refers to the solution for connecting VoIP traffic and data traffic between sites over an IP network capable of providing guaranteed-bandwidth services.

Voice-Trunking and Toll-Bypass Strategy

In this scenario, the customer transports the voice traffic on traffic-engineering tunnels. The VoIP bandwidth requirements between two sites are admitted on a guaranteed-bandwidth tunnel. The guaranteed-bandwidth tunnel uses a dedicated-bandwidth pool, either the main pool or the subpool.

On the provider edge, the VoIP traffic is forwarded on the traffic-engineering tunnel in two methods: static routes and policy-based routing. Using static routes to forward the VoIP traffic is possible when there is a specific destination subnet or destination hosts receiving the VoIP traffic. If the same destination is supposed to receive VoIP and data traffic, PBR can be used to distinguish the VoIP traffic and forward it appropriately on the traffic-engineering tunnel. With PBR, the VoIP traffic can be identified with any of the extended ACL parameters, such as the Real-Time Transport Protocol (RTP) port number, or premarked type-of-service (ToS) bits of the voice packets. Because the VoIP trunking could have multiple destinations, the matching should be on the destination prefix and either the ToS bits or the destination port numbers. To identify the VoIP traffic based on the ToS bits, the ToS bits need to be set in advance at the customer edge using various marking techniques.

In order to make sure that the traffic sent on a guaranteed-bandwidth tunnel is conforming to the bandwidth reserved, additional traffic shaping and rate limiting is required on the customer edge.

Note: This example assumes that the VoIP gateway functionality is performed on the gateway located at the customer-edge site. This example covers voice connection trunk configuration, which creates a permanent point-to-point voice circuit from a provider edge. For details on configuring voice connection trunk go to: http://www.cisco.com/warp/public/788/signalling/trunk_config.html.

Note: In addition, at the time this document was created, PBR is not supported for VPN traffic routed from a VRF interface.

Voice-Trunking and Toll-Bypass Network Topology

Figure 7 shows the toll-bypass network diagram for the central-office or private branch exchange (PBX) with a packet interface.

Figure 8 shows the voice-trunking network diagram for a VoIP-over-packet network.

Figure 7
Network Diagram for Toll Bypass over Voice-Enabled Network

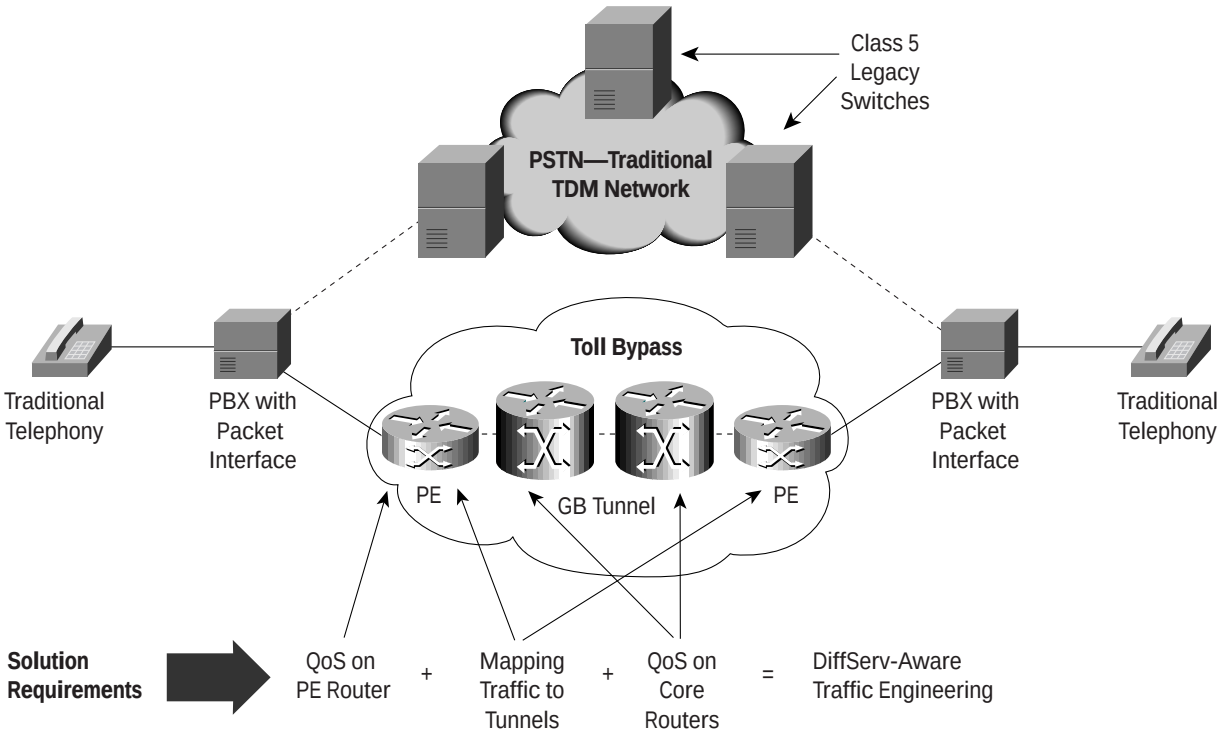

Figure 8
Network Diagram for Toll Bypass with Voice/Data-Converged Network

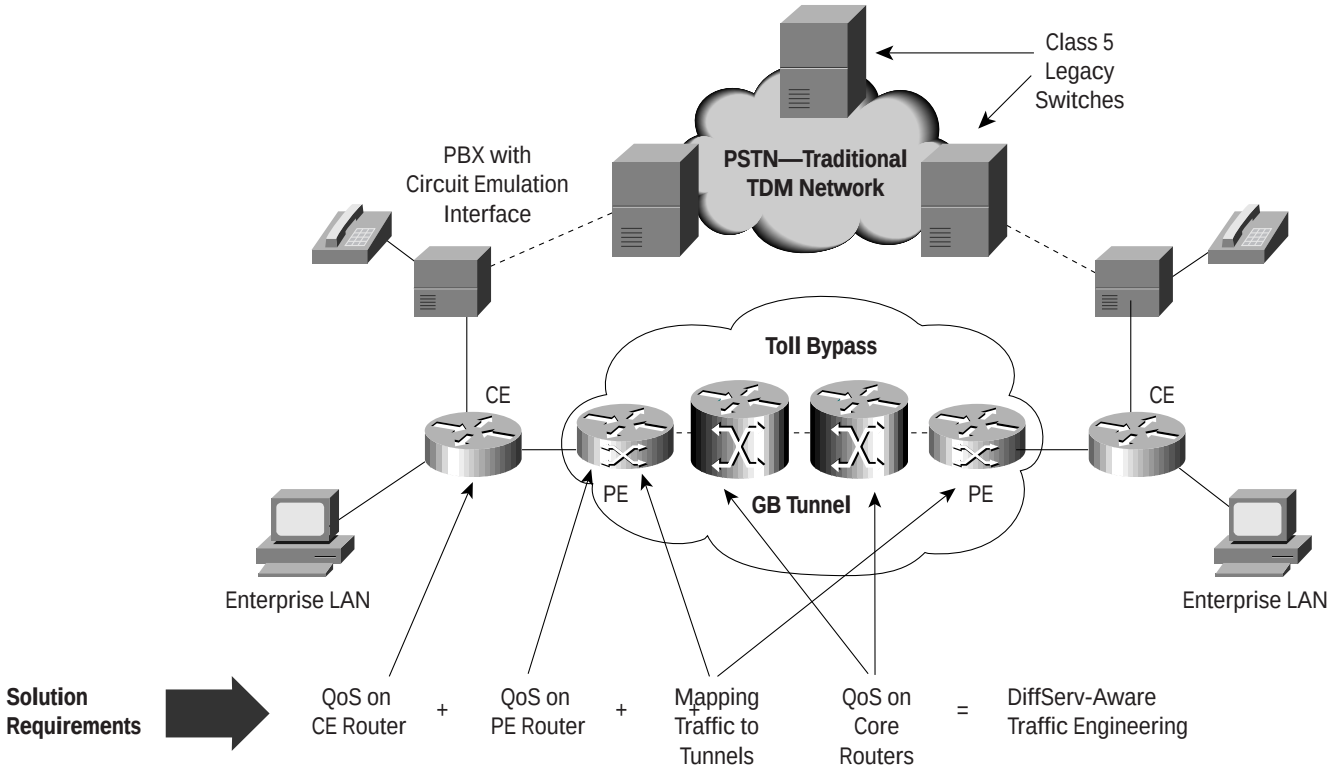

Voice-Trunking and Toll-Bypass Benefits

The benefits of deploying the VoIP trunking scenario are as follows:

- Voice trunking and toll bypass offer new premium services for high-priority traffic, such as voice traffic with tight guarantees for throughput, delay, and reliability.
- Voice trunking and toll bypass provide point-to-point bandwidth admission and guarantees of the VoIP traffic.
- Voice trunking and toll bypass reduce costs by taking advantage of available bandwidth; they also provide guarantees for high-priority traffic.
- Voice trunking and toll bypass offer the ability to provision voice capacity on the network and manage core links utilization.
- Voice trunking and toll bypass provide high availability for VoIP traffic with FRR.

Voice-Trunking and Toll -Bypass Ramifications

The ramifications of deploying this VoIP trunking scenario are as follows:

- The voice-trunking and toll-bypass network is more complex to implement and troubleshoot than an IP network.
- Policing of the VoIP traffic at provider edge with MQC is for the aggregate of all the VoIP flows—not for each flow.
- Extensive provisioning and automation tools are necessary to enable large-scale deployment.

Voice-Trunking and Toll-Bypass Configuration Summary

The following is a summary of tasks to configure the edge router for voice trunking and toll bypass:

Step 1. Admit the VoIP trunking onto the traffic-engineering bandwidth pool.

The following configuration reserves 256 kbps into the subpool bandwidth.

```
Interface tunnel 12
 Tunnel mpls traffic-eng bandwidth sub-pool 256
```

In this example, the traffic-engineering tunnel headend is configured with a traffic-engineering tunnel for the voice trunk and is using the subpool. Refer to the "Create Traffic Engineering Tunnel Interface" step for a discussion on traffic-engineering headend location alternatives.

Step 2. Inbound VoIP traffic at the ingress provider edge.

The following configuration uses MQC to identify the incoming IP traffic with Precedence 5 as the VoIP class and polices it to 256 kbps, setting the MPLS EXP bits to the guaranteed-bandwidth services Class 5. This configuration assumes that the customer-edge router did set the VoIP packet precedence bits to 5.

```
!
class-map match-all VOIPTOSItraffic engineering2
 match ip precedence 5
!
policy-map IN-POLICY
 class VOIPTOSItraffic engineering2
 police 256000 48000 96000
 conform-action set-mpls-exp-transmit 5
 exceed-action drop
 !
 interface Serial1/0
 service-policy input IN-POLICY
!
```

The bandwidth policed on the inbound interface of the ingress provider edge is the aggregate of all the point-to-point voice traffic sent from this site.

Note: Additional methods, such as matching and marking the traffic, are not discussed in this document. For more information on Configuring QoS Features go to h[illegible] Also, after the traffic has entered the ingress provider edge and has being marked, the outbound physical interface on the provider edge and every hop along the path will treat the traffic based on its marking and the interface configuration shown in Step 2 in the section "Guaranteed-Bandwidth Services Core Configuration Summary."

Step 3. Identify outbound VoIP traffic at the egress provider edge.

In this step, configure low-latency queuing for the voice-trunking class on the outbound interface of the egress provider edge. In the following configuration, 20 percent of the link bandwidth is reserved for all the voice-trunking traffic with ToS 5 set, and received from all the guaranteed-bandwidth services tunnels.

```
!
Class-map match-all VOIP
 match ip precedence 5
!
policy-map OUT-POLICY
 class VOIP
  priority percent 20
!
```

```
Interface Serial2/0
 service-policy output OUT-POLICY
!
```

In this configuration, the VoIP traffic is receiving priority queuing, with 20 percent of link speed with MQC. The bandwidth reserved for the priority queue should be enough for all the VoIP trunks terminating at the site. For additional information on configuring MQC go to: [illegible]

Step 4. Forward the traffic onto the tunnel.

The following configuration at the tunnel headend router forwards the voice-trunking traffic onto the provisioned traffic-engineering tunnel using static route.

```
ip route 10.20.35.1 255.255.255.255 tunnel12
```

In this example, the static-route command forwards voice traffic to a dedicated VoIP gateway. In cases where a site will be sending VoIP and data to the same destination, IP PBR with extended ACLs can be used to distinguish between types of traffic, and forward it appropriately.

Step 5. Verify the configuration.

Use the following command to verify the QoS configuration:

```
mq-7500c#show policy-map
 Policy Map OUT-POLICY
 Class BUSINESS
 bandwidth percent 50
 Class VOIP
 priority percent 20
 Class class-default
 fair-queue
 random-detect

mq-7500c#show policy-map interface serial 2/0
 Serial2/0
 Service-policy output: OUT-POLICY (1201)
 queue stats for all priority classes:
 queue size 0, queue limit 2210
 packets output 0, packet drops 0
 tail/random drops 0, no buffer drops 0, other drops 0
 Class-map: VOIP (match-all) (1207/5)
 34523 packets, 1234564 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
 Match: ip precedence 5 (1209)
 Priority: 20% (8842 kbps), burst bytes 221050, b/w exceed drops: 0
 Class-map: class-default (match-any) (1211/0)
 0 packets, 0 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
 Match: any (1213)
 0 packets, 0 bytes
 5 minute rate 0 bps
 queue size 0, queue limit 3315
 packets output 0, packet drops 0
 tail/random drops 0, no buffer drops 0, other drops 0
 Fair-queue: per-flow queue limit 828
 Random-detect:
 Exp-weight-constant: 9 (1/512)
 Mean queue depth: 0
 Class Random Tail Minimum Maximum Mark Output
```

```
drop drop threshold threshold probability packets
0 0 0 828 1657 1/10 0
1 0 0 931 1657 1/10 0
2 0 0 1035 1657 1/10 0
3 0 0 1138 1657 1/10 0
4 0 0 1242 1657 1/10 0
5 0 0 1345 1657 1/10 0
6 0 0 1449 1657 1/10 0
7 0 0 1552 1657 1/10 0
```

Related Documents

(1) MPLS Traffic Engineering

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120s/120s5/mpls_te.htm

(2) DiffServ-Aware Traffic Engineering

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120st/120st11/ds_te.htm

(3) RFC 2702—Requirements for Traffic Engineering over MPLS

(4) RFC 2205—Resource Reservation Protocol (RSVP)

(5) draft-katz-yeung-ospf-traffic-03.txt

(6) MPLS Traffic Engineering Fast Reroute — Link Protection

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120st/120st16/frr.htm

(7) Policy-Based Routing

http://www.cisco.com/warp/public/cc/techno/protocol/tech/plicy_wp.htm

(8) MPLS Class of Service

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/cos.htm

(9) Configuring MDRR/WRED on a Cisco 12000 Series Router

(10) Quality of Service Command-Line Interface Overview

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt8/qcfmdcli.htm

(11) Class-Based Weighted Fair Queuing

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/cbwfq.htm

(12) Quality of Service for Voice over IP

http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/qossol/qosvoip.htm

(13) Voice: Configuring Connection Trunk

(14) draft-ietf-mpls-rsvp-lsp-tunnel-08.txt: RSVP-Traffic Engineering Extensions to RSVP for LSP Tunnels

(15) Configuring the Modular Quality of Service Command-Line Interface

(16) IP Routing Protocols Configuration

(17) Cisco IOS Quality of Service Solutions Configuration Guide

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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 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
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

xxxxxx/ETMG 9/02